THE UNIVERSITY
of EDINBURGH

This thesis has been submitted in fulfilment of the requirements for a postgraduate degree (e.g. PhD, MPhil, DClinPsychol) at the University of Edinburgh. Please note the following terms and conditions of use:

ESSAYS ON DYNAMIC PROCUREMENT

By

JUSTUS LAUGWITZ

A dissertation submitted in partial fulfilment of
the requirements for the degree of

DOCTOR OF PHILOSOPHY

UNIVERSITY OF EDINBURGH
School of Economics

JANUARY 2020

DECLARATION

I declare that this thesis has been composed solely by myself and that it has not been submitted, in whole or in part, in any previous application for a degree. Except where stated otherwise by reference or acknowledgement, the work presented is entirely my own.

Justus Laugwitz   Edinburgh, January 2020

ACKNOWLEDGEMENT

ESSAYS ON DYNAMIC PROCUREMENT

Abstract

by Justus Laugwitz,
University of Edinburgh
January 2020

This thesis studies the impact of uncertainty and its sequential resolution on the design of procurement mechanisms and, in particular, auctions. Procurement contracts constitute a considerable fraction of output in developed economies. These contracts are typically long-term agreements between a buyer (e.g., the government) and a seller (a supplier of a good or service). This long-term nature of procurement contracts makes renegotiation after the resolution of uncertainty more likely. As a consequence, analysing how these ex-post dynamics affect the ex-ante design of procurement mechanisms becomes very important.

The first chapter of this thesis provides a survey and critical review of the theoretical literature on dynamic procurement. It provides a summary of the seminal contributions studying the optimal design of procurement mechanisms and reviews them in light of recent advancements in the theory of dynamic procurement. We describe the issue of procurement from a mechanism design standpoint, the dominant scoring auction, as well as the problems of adverse selection and corruption. The second part of this review describes different channels through which dynamic considerations impact this environment. In particular, we consider issues such as renegotiation and ex-post adaptions, the non-contractible nature of some relationships, repeated interactions, as well as the unique aspects relevant for large and

complex projects. In the final part of the first chapter, we explore how the findings from the literature on dynamic procurement help us cast new light on the key static ideas from the seminal papers.

The second chapter of this thesis explores the effect of cost-uncertainty on the optimal procurement mechanism. We make three main contributions to the literatures on mechanism design and procurement. First, contributing to the research on dynamic mechanism design, we derive the optimal mechanism for a multidimensional environment with sequential generation of private information. Unlike in the one-dimensional analysis of Eső and Szentes (2007), in our multidimensional dynamic model, the second stage private information does affect the agents' rent. Second, contributing to the literature on procurement auctions, we show that, unlike in Che (1993), the scoring auction is no longer able to implement the optimal mechanism if contracts are renegotiated in response to the realisations of cost shocks. Finally, the paper introduces a stylised model of procurement with renegotiation that matches several features of observed procurement auctions.

The third and final chapter of this thesis addresses a question that naturally emerges from the results of the second chapter. Scoring auctions with renegotiation are not the optimal procurement mechanism for the buyer. However, despite their sub-optimality, scoring auctions are the predominant mechanism used in practice. This chapter addresses the question of optimality within the class of scoring auctions: What is the optimal scoring rule? Furthermore, we quantify the gap between the overall optimal mechanism and the optimal scoring auction in terms of buyer payoffs. In addressing these questions, this chapter introduces a methodology that imposes further restrictions on a direct mechanism optimisation problem to find the optimum among a class of sub-optimal mechanisms. We further extend the result of the second chapter to a discrete type-space. This type-space facilitates an analytic solution to the optimisation problem for the restricted mechanism design problem. Finally, we introduce an auction with reserve-scores and show that it is the optimal scoring auction with renegotiation for discrete type-spaces.

LAY SUMMARY

This thesis studies how a buyer, such as a government, should design the process through which she purchases goods and services when many sellers are competing to provide it. Several distinct features often characterise the procurement process. Usually the good or service is not offered as a one-off, but rather over a substantial length of time. As a consequence procurement contracts are commonly long-term agreements. Moreover, in procurement, the good or service to be purchased is not usually valued exclusively in terms of the price. For example, if a government decides to build a road, it is not just concerned about the cost of the road, but considerations such as the length of disruption to traffic or the expected time until repairs are necessary are also important. We refer to this aspect of procurement as multidimensionality. Finally, as a consequence of the multidimensionality and long-term nature of procurement contracts, they are frequently renegotiated in practice. I.e. after the initial agreement between the buyer and one or more sellers, the parties may alter or add to the purchasing contract.

In this thesis, we address how the dynamics arising from renegotiation should affect the initial process through which the buyer allocates the good to a seller. In the first chapter, we review existing literature on this topic. Here we summarise recent developments in the literature and explore the different approaches to modelling a dynamic procurement process. In the second chapter, we introduce a model of procurement that is both multidimensional and allows for the possibility of renegotiation. With this model, we describe the properties, such as expected payment by each seller, an optimal auction would have to satisfy. However, in this chapter, we show that there is a strong negative result. One of the most commonly used auctions for procurement, the scoring auction, is not optimal for the buyer. Finally, in

the third chapter, we introduce a simplified version of the model from chapter 2. This model allows us to explore the optimal mechanism in more detail. Furthermore, it enables us to describe how much is lost relative to the optimum if we do use the scoring auction that is popular in practice.

TABLE OF CONTENTS

LIST OF FIGURES

Chapter One

DYNAMIC ISSUES IN THE THEORY OF PROCUREMENT

1.1   Introduction

Procurement, both public and private, makes up a considerable fraction of total economic activity in developed nations. According to the handbook of procurement, not accounting for the value of transactions and private sector procurement, public procurement alone made up 16% of GDP in EU countries and 20% in the US in 2009 (Dimitri, Piga, and Spagnolo, 2006). Recent estimates suggest this number is steadily increasing. As a consequence, procurement has received considerable attention in the economic literature. This literature is in part reviewed by Klemperer (1999) and Armstrong and Sappington (2007) with a focus on procurement auctions and the theory of regulation, respectively. Key issues in the literature include the design of the optimal procurement auction and the impact and prevention of corruption, among others. In this review, we will summarise the literature of procurement while focusing on one particular aspect of it: the dynamics arising from the long-term nature of procurement contracts.

Many procurement contracts are long-term agreements. For example, contracts for the provision of public utilities, construction projects, or supply agreements for complex goods

such as machinery can last anywhere from multiple years up to several decades. In these kinds of long-term commitments, ex-post adaptations of the initial agreement are commonplace. Evidence for this is readily available. E.g., Arve and Martimort (2016) cite evidence that most of UK procurement contracts for public utility provision are revised ex-post. Bajari, Tadelis, et al. (2001) argue that procured construction projects frequently use change orders to adapt the initial agreement whenever needed. For public procurement projects, Jung et al. (2019) claim that most, if not all, projects are renegotiated after the initial tender.

Given the likelihood of ex-post adaptation for the initial agreement in procurement, we might wonder what impact this has on the procurement market as a whole. Contributions in the theory of procurement have considered many different channels through which the dynamics affect the strategic interaction between sellers and the buyer in procurement. Early contributions such as by Dewatripont (1989) and Laffont and Tirole (1990) considered how the possibility of adjusting contracts ex-post changes the incentives in the initial allocation setting. We can show that long-term procurement contracts are not 'renegotiation proof'. 'Renegotiation' here simply refers to any adjustment of the initial contract without necessarily having to specify the mechanism of renegotiation. 'Renegotiation proof' thus means the initial contract is set up in such a way that renegotiation is not incentive compatible for both sides, i.e. mutual gains from re-agreeing the contract cannot be made. As a consequence, incentive distortions that may be favourable to the buyer in the allocation may not be possible due to the possibility of renegotiation. In response to this finding, some, such as Wang (2000) and Bajari, Tadelis, et al. (2001), considered whether the allocation stage of long-term procurement contracts is meaningful at all or just 'cheap talk' on the way to ex-post adaptations. Others, such as Arve and Martimort (2016) or the later chapters of this PhD thesis attempted to consolidate the findings of incentive theory and mechanism design with dynamic models of procurement. We will introduce these and other contributions in section 1.3 of this paper.

We might further wonder how if at all, the key findings of seminal static procurement

analyses would change if we account for the fact that the initial assumption of a static procurement environment is so often not satisfied. For example, is the optimal static multidimensional mechanism of McAfee and J. McMillan (1987), Riordan and Sappington (1987a), and Laffont and Tirole (1987) used in much of the procurement literature still optimal when there is initial uncertainty that is resolved before the contract is (fully) implemented? Does the Che (1993) prediction that scoring auctions can implement the optimal procurement mechanism hold up in the face of this? Or how are the results on the design of the optimal procurement auction by Branco (1997), Asker and Cantillon (2005), and Asker and Cantillon (2008) impacted by allowing for potential dynamics? Are the predictions for corrupt procurement auctions of Ganuza (2007) and Burguet and Che (2004) affected by these dynamics? We will describe a selection of the most influential findings from the part of procurement theory that treats the environment as static and review them in light of the results of the dynamic literature.

This paper is set up as follows. Section 1.2 will introduce a selection of seminal works on static procurement to highlight key issues that have been addressed in the literature. Section 1.3 will follow up by giving a summary of recent developments in the theory of dynamic procurement. Section 1.4 will then discuss how the findings of section 1.3 can shed light on the results of the papers from section 1.2 from a different perspective, accounting for the possibility of dynamic effects. Section 1.5 will conclude and summarise the discussion presented.

1.2   Seminal Contributions to the Theory of Procurement

A defining feature of procurement auctions is that they usually do not meet an assumption central to most other auction environments. The assumption here is that the good to be auctioned is a well-defined object, and thus the only aspect to be determined is the price. In most procurement environments, this assumption is not tenable. For example,in construc-

tion, the build time and design specification of a project may be of considerable importance and receive appropriate attention in the tender. Even seemingly straightforward goods, such as latex gloves for medical services, vary considerably across firms in design, reliability, material, speed of delivery, and many other factors the auctioneer (buyer) may have preferences over. In this multidimensional context, especially if firms face varying costs of production for the non-price elements, the optimal allocation mechanism should take these into account. In the literature on multidimensional procurement, for simplicity, most papers focus on one composite non-price aspect we will refer to as 'quality'. This is in line with the notation of Che (1993).[1] Quality might refer to any non-price factor relevant to the procurer. Restricting attention to this composite element allows us to develop stylised models which are still able to capture many of the crucial issues arising in multidimensional environments.

McAfee and J. McMillan (1987), Riordan and Sappington (1987a), and Laffont and Tirole (1987) initially describe the optimal allocation mechanism relevant for the multidimensional procurement context. Here the trade-offs for the procurer are considerably different compared to single dimensional environments. Unlike the latter, if the model is multidimensional, the procurer faces two potential inefficiencies – allocative inefficiency and quality inefficiency. The former arises when the good is not allocated to the seller with the lowest cost of producing quality or not allocated to any seller. In this situation, we can achieve a higher social surplus by changing the allocation. In other words, the allocation is inefficient. The latter inefficiency arises when the seller produces the good at a quality that does not maximise the difference between benefit (for the buyer) and cost (for the seller). Here the quality is inefficient since changing it could increase social surplus. It is a well-known result in single-dimensional auction environments that the auctioneer should sacrifice allocative efficiency to reduce the informational rent obtained by the bidder (Myerson, 1981). In multi-dimensional

---

[1]We can equivalently think of this as 'effort' such as in Laffont and Tirole (1987), or the overall design specification of the procurement project. Depending on the exact contex, either definition might be more appropriate.

4

models, this trade-off is not usually present. Instead, the buyer-optimal mechanisms of McAfee and J. McMillan (1987), Riordan and Sappington (1987a), and Laffont and Tirole (1987) solely introduce inefficiency in the quality of the good to achieve a similar reduction in the sellers' rent. A trade-off to reduce the quality aspect of the procurement good below its social optimum allows the buyer to extract a higher surplus by reducing the surplus of all other agents.

One of the most predominantly used mechanisms in procurement, when the good is multidimensional, is the scoring auction. This mechanism allows the bidders to bid a design specification, achieving scores on different non-price elements according to the scoring rule set up by the auctioneer. For example, an auctioneer might decide to reward an earlier expected completion date for a building project in the scoring rule. Che (1993) describes different scoring auctions for two-dimensional procurement, where bidders bid a price as well as a composite quality level. The auctions described by Che closely match the most common single-dimensional auction mechanisms. For instance, in a sealed-bid first-score auction the bidders privately submit a bid consisting of a price p and a quality q that achieve score $S(p,q)$ according to the scoring rule S laid out by the auctioneer. The contract is then awarded to the bidder who bid the highest score. He will be asked to implement q and paid p. In intuition and result, this mechanism closely resembles the first-price auction from price-only auction theory. Che shows that if the buyer can commit to a scoring rule and if the sellers' types are independent of each other, the scoring auction can implement the optimal mechanism of McAfee and J. McMillan (1987), Riordan and Sappington (1987a), and Laffont and Tirole (1987) with a quasilinear scoring rule.

In these analyses, although the good is treated as multidimensional, the assumption is maintained that sellers' private information only varies in one dimension: in their cost of producing quality.[2] Asker and Cantillon (2008) and (2010) argue that this assumption is unrealistic. They extend the McAfee and J. McMillan (1987), Riordan and Sappington (1987a),

---

[2]Or exerting effort.

Laffont and Tirole (1987), and Che (1993) papers to allow not just for a multidimensional object but also multidimensional private information. This extension enables the authors to better match the equilibrium behaviour from the scoring auction model to what is observed in the data (Asker and Cantillon, 2008). They are also able to derive the optimal mechanism for this environment. However, although the scoring rule gets closer to implementing it than a simple bargaining mechanism, a scoring auction is no longer optimal (Asker and Cantillon, 2010).

For simplicity and stylisation, many multidimensional models of procurement restrict attention to environments where the sellers' private information is not correlated. This is a strong assumption that arguably does not correspond well with reality in many procurement settings. For example, in construction, many projects require subcontracting and firms may choose among a similar pool of firms to subcontract to. Such firms' costs are unlikely independent of each other. Branco (1997) agrees with the papers described above that one key aspect of procurement auctions is the complexity of the strategy space, where firms compete in several dimensions over the allocation of the contract. However, he argues that the assumption of independence between the firms' cost types maintained by Che (1993) is unrealistic in the procurement environment. Branco shows that if an auction is both multidimensional and has bidders with correlated types, the design of the optimal auction is very different from auctions with either homogeneous goods or independent types. In particular, the optimal auction must consist of two stages, wherein the first stage a bidder is selected who negotiates the quality aspect with the auctioneer in the second stage. Moreover, unlike in Che (1993), the quality will be a function of all bidders' bid qualities and as such a two-dimensional scoring auction is not an appropriate implementation mechanism.

Manelli and Vincent (1995) look at the quality aspect of procurement auctions through a different lens than the papers introduced so far. Here quality is not something chosen by the sellers. Rather each seller is endowed with the ability to produce the good at a fixed level of quality. Moreover, this quality is the private information of each seller. The sellers know

the quality of their good and the buyer's valuation of the good is dependent on the selected seller's good's quality. This model lends itself to explain the potential for adverse selection in procurement auctions. If quality is not ex-post verifiable, an auction mechanism such as a first-price auction may not select the seller who would produce the greatest economic surplus in many environments. A seller with a low quality good may be able to produce the good more cheaply and will thus be able to undercut the high-quality sellers in an auction. This finding is contradictory to the result of Myerson (1981) that auctions, as they induce competition among bidders, reach a socially desirable outcome. In the Manelli and Vincent (1995) analysis, the buyer may be better off selecting sellers individually and giving them a take-it-or-leave-it offer. Manelli and Vincent describe the exact conditions for when the adverse selection effect makes it such that this process outperforms the simple first-price auction.

A final aspect that has received considerable attention in the procurement literature is the issue of corruption. Procurement is often seen as susceptible to corruption in both the media as well as the economic literature. Much EU as well as government policy is dedicated to reducing the presence of corruption in procurement. Two critical questions on corruption in procurement are at the heart of most of the literature. The first is on the source of the corruption. One argument here is that it exists due to a lack of competition. When firms can extract large rents it may create a strong incentive for them to appropriate these (Rose-Ackerman, 1996). Celentani and Ganuza (2002) describe a model of procurement with equilibrium corruption that allows for a theoretical analysis of this claim and tests whether increased competition among sellers would lead to a decrease in corruption. Here the authors find that the Rose-Ackerman (1996) prediction does not generally hold: competition among sellers has an ambiguous effect on corruption.

The second key question addressed in the literature on corruption in procurement is on how to limit its adverse impact. In an early contribution, Laffont and Tirole (1991) discuss measures to achieve reductions in corruption for procurement auctions. They find that when

there is a possibility of collusive behaviour, auction mechanisms that do not distinguish between non-price elements tend to perform better. In effect, corruption limits the possibility of bidding on quality. This result is also found by Burguet and Che (2004) who endogenise the favouritism by introducing a corruptible third-party responsible for evaluating the quality-bids of sellers. Since the corruption occurs through the channel of non-price elements, the optimal (scoring) auction in this environment de-emphasises the importance of quality. This model additionally shows two surprising results. First, if the power of the third-party is limited, the corruption can benefit the buyer, since it reduces the informational rent of the agents by forcing them to bid more aggressively. Second, the optimal corruption-limiting auction may have favouritism embedded. By lowering the handicap on the efficient firm, the auctioneer increases this firm's bidding advantage and limits the scope for the other firm to use corruption.

1.3   The Literature on Dynamic Procurement

The literature so far discussed has mostly treated the procurement problem as static. Many sellers are willing to provide a good or service to the single buyer and once agreed, the contract is assumed to play out as specified. In reality, this assumption is hardly tenable. The frequency of ex-post renegotiation of procurement contracts is well documented. On top of this, repeated interaction and consequently reputation effects are commonplace. Although for example, the model of Branco (1997) features a dynamic implementation, this is still with a static environment in mind. In this section, we will introduce papers that explore how the dynamic nature of the procurement environment should impact our understanding of it.

In the wider auction literature, the optimal design of auctions and mechanisms for dynamic settings is a broad and active literature. Bergemann and Said (2010) classify dynamic auctions into two broad sets: (i) auctions with a dynamic population of agents with fixed private information, and (ii) auctions with a fixed population of agents whose private infor-

mation changes across time. Bergemann and Said summarise the literature for both these subsets of dynamic auctions. An example of (i) is the analysis of optimal dynamic auctions with index rules by Pai and Vohra (2013). Here many agents arrive at the auction over time and may choose to submit a bid for one of several goods before a deadline. This environment is similar to the sale of flight tickets online and Pai and Vohra (2013) show that an index rule, similar in concept to the scoring rule of Che (1993) for the procurement setting, can implement the optimal allocation under standard assumptions on the distribution of valuations. A recent example of (ii) is the theory of indicative bidding by Quint and Hendricks (2018). Selling a business through an auction is one of the most common ways auctions are used in practice in terms of monetary value sold. In these environments nonbinding preliminary (indicative) bids are a commonly used mechanism to filter out serious bidders from a large pool. Quint and Hendricks (2018) describe an auction mechanism with indicative bids formally and show that it outperforms an open auction when entry is costly in many environments. This paper fits into (ii) since agents learn their true valuation only after submitting an indicative bid. Unlike the literature review by Bergemann and Said (2010), in our review we restrict attention to specifically the field of procurement.

1.3.1 Procurement and Renegotiation

Many possible reasons for the frequency of renegotiation in the procurement environment have been explored in the literature. The long-term and multidimensional nature of procurement contracts makes it challenging to write a complete contract for every contingency that may arise over the course of it. Other factors, such as poor governance and regulatory weakness, can also play a role. In the theory of renegotiation of procurement, it is often a different impetus for renegotiation that leads to the varying conclusions of the models. Estache et al. (2009) provide a good introductory summary of the issue of renegotiation in multidimensional procurement. In this part of the literature review, we will summarise the

key findings of a selection of papers on renegotiation in light of the varying reason for why renegotiation occurs. We will see that this reason is strongly connected to the conclusions of the papers and helps to explain the difference in their findings.

In an early contribution to the literature on renegotiation in procurement by Tirole (1986), renegotiation arises due to incomplete contracts. This model predates the optimal mechanism for multidimensional procurement. Instead, a two-stage model is introduced. Here a seller invests in a project and then he and the buyer realise their cost and value of the project afterwards. Tirole examines the question of whether we should expect underinvestment in this context. Without having to specify a bargaining mechanism, Tirole shows that underinvestment will occur so long as the bargaining process is inefficient. I.e., as long as the seller loses some rent from the investment he will underinvest in equilibrium. This underinvestment relative to the social optimum may be considerable depending on the bargaining mechanism.

The impetus for renegotiation in Laffont and Tirole (1990) comes from the seminal work of McAfee and J. McMillan (1987), Riordan and Sappington (1987a), and Laffont and Tirole (1987) on the optimal mechanism for multidimensional procurement environments. In particular, Laffont and Tirole (1990) examine what would happen if we allow for the distorted optimal qualities of the multidimensional procurement mechanism to, later on, be renegotiated. They show that generally for any previously agreed quality, the principal would like to ex-post reagree the quality with the agent. By moving from distorted to the socially optimal qualities, gains from trade can always be made. If the principal cannot commit not to renegotiate the qualities in this way, however, it becomes impossible for her to implement any incentive-distorting qualities from McAfee and J. McMillan (1987), Riordan and Sappington (1987a), and Laffont and Tirole (1987). The qualities of the first stage become 'cheap talk' waiting to be renegotiated.

A critical question that arises when we allow in our models of procurement for the possibility of renegotiation is the question of when the bids for the project in the allocation

phase are meaningful or just 'cheap talk'. In Tirole (1986) and Laffont and Tirole (1990), the first stage of the procurement mechanism was cheap talk as a consequence of a lack of commitment. Wang (2000) addresses this question in a model of auctions with renegotiation that allows for commitment. In this model, the impetus for renegotiation comes from the commitment rather than from a lack of it. Here the buyer can commit to not awarding the contract to any of the bidders. This commitment power allows the government to use potential information it might learn about the bidders in the auction to approach them with take it or leave it offers rather than accept their offers. Wang shows that in this model, the initial bids matter and are not just 'cheap talk'. Intuitively, since renegotiation is not guaranteed to take place, when it does take place the buyer can learn about the sellers' cost from their bids.

The Wang (2000) model gives one reason for why renegotiation might occur in practice, namely that it can help the buyer obtain lower prices. Arguably, however, in practice, this is not usually the reason why procurement contracts are renegotiated. Instead they tend to be renegotiated either because changing a design specification can be beneficial to the buyer or the seller or, as Chang, Salmon, and Saral (2016) argue, because the seller would otherwise go bankrupt. Chang, Salmon, and Saral describe the procurement environment as one that (i) very commonly has a common value aspect to the auction and (ii) as one whereas a consequence the winning bidder frequently suffers from the winners curse. Furthermore, in equilibrium, the winner's curse should be eliminated, so observing it in practice points to something else going on. In their 2016 model, they allow for wealth constraints to examine whether renegotiation can be an effective tool in mitigating the winner's curse. Generally, this is not the case, however, except for those smaller sellers who can credibly threaten to go bankrupt. In equilibrium, renegotiation thus mostly benefits weaker sellers by allowing them to make use of this commitment device to submit a more competitive bid. This model therefore suggests that smaller firms may use the commitment to go bankrupt to force beneficial renegotiation in the procurement process.

In an essential contribution to understanding long term dynamics in procurement, Arve and Martimort (2016) focus on the effect of risk aversion and cost observability on the design of the optimal dynamic contract. They consider a two-period model of a procurement process that involves an uncertain add-on. Arve and Martimort consider both the cases where the cost of the second-period add-on is verifiable and where it is the seller's private information. In either case, the necessity of the add-on can be expected, and the buyer and seller are allowed to contract over it with full commitment in the initial period. Unlike in the former case, in the more realistic second scenario where the cost is private information, Arve and Martimort show that under the assumption of risk aversion the rent from this private information will be distorted through the insurance channel.

Arve and Martimort show that allowing for risk aversion changes the incentives present for the firm in both periods of the model. They distinguish between two main effects. The 'income effect' of risk aversion refers to reduced first-period distortions of the optimal contract from moving payments to the second stage and in this way reducing the informational rent of the firm. On the other hand, the 'risk effect' refers to a reduction of the size of the second stage add-on below the first-best efficient level stemming from the firm having to bear part of the add-on risk in order to reveal its costs in the second stage. This risk cannot be fully insured against it.

The literature introduced so far described several different channels of how renegotiation can occur in procurement. Bajari, Tadelis, et al. (2001) go one step further and argue that "[...] the procurement problem is primarily one of ex-post adaptations rather than ex-ante screening". Their 2001 paper argues that previous literature is too fixated on the issue of ex-ante screening and evidence from management science shows that this focus is misplaced as little evidence points to the existence and relevance of strong ex-ante informational discrepancies. As such, in their model Bajari, Tadelis, et al. take the opposite approach and completely ignore the ex-ante hidden information problem. In this way, they can restrict attention to the adaptation to changes arising from the inherently incomplete design. In

this model, the sellers have private information about the cost of these changes and Bajari, Tadelis, et al. detail a trade-off between effort inducing incentives and adaptability to changes in the product design. The more complex the procurement project, the more beneficial become incomplete ('cost-plus') contracts with high adaptability. On the contrary, simple projects are best served by complete effort inducing ('fixed-price') contracts.

Strategic Bidding

A key consideration for models of renegotiation in procurement is how the presence of renegotiation will affect the strategic bidding in the allocation stage of the model. The models we introduced in the previous section address this to a varying extent. However, it is worth examining those papers that explicitly focus on these issues.

There is strong evidence that firms respond to the presence of renegotiation in their bidding behaviour. In their (2014) paper Bajari, Houghton, and Tadelis develop an empirical model to test for the presence and implications of adaptation cost to renegotiation in procurement. The ask if there is evidence for strategic bidding in response to the incompleteness of contracts. According to their estimation, adaptation costs are shown to be considerable, and there is support for the existence of strategic bidding in response to it. They find evidence in support of a Bajari, Tadelis, et al. (2001) thesis. Namely, when sellers expect renegotiation in response to incompleteness, it should reduce their incentive to bid for this contract if renegotiation is costly. Furthermore, this effect takes a higher priority for the sellers' bidding behaviour in the allocation mechanism than the initial private information.

Jung et al. (2019) extend the Bajari, Houghton, and Tadelis (2014) analysis to incorporate bid-level data. This additional information allows the authors to obtain an estimate of the extent to which firms shade their bids in a procurement auction, i.e. their projected markup over cost for each item. With this analysis Jung et al. find a positive correlation between the bid markup and the presence of ex-post (quantity) renegotiation. Firms conversely lower the markup on non-renegotiated items to maximise their surplus after the allocation of the

contract. This behaviour leads to increases in cost to the procurer that may far exceed previous estimates. Given how commonplace renegotiation is in procurement – Jung et al. (2019) argue that most, if not all, procurement contracts are later renegotiated – possibly too much of the procurement literature is reduced to a static private information problem rather than considering its dynamic issues. Failing to account for the considerable effects the frequent post-tender renegotiation of procurement contracts has on the design of the contract and the nature of the optimal allocation mechanism could lead us to make significant errors in predicting the social and buyer-optimal mechanisms for this context.

1.3.2 Cost Overruns and Complex Projects

One aspect of procurement that receives a particularly large amount of media coverage is the occurrence of significant cost overruns in sizeable public procurement projects. These cost overruns are often due to changes in the initial design specification that were not anticipated at the time of the allocation of the contract. Ganuza (2007) discusses the problem of cost overruns in complex procurement contexts as a problem of underinvestment. In his model, a procurer faces a trade-off to invest in the initial design specification. Higher investment reduces the likelihood of costly renegotiation and increases the probability of selecting the most efficient firm for the project. However, the lower investment allows the procurer to reduce the competitive advantage of efficient firms and thus reduces informational rent. Ganuza's model predicts underinvestment in the design specification as a result of this trade-off – leading to extensive renegotiation.

Herweg and Schmidt (2017,2018) and Herweg and Schwarz (2018) approach the feature of large, complex procurement projects to have significant cost overruns from a different perspective. Herweg and Schmidt (2018) consider a model of large procurement projects in which the buyer is aware of potential flaws in her design. Herweg and Schmidt argue for an arbitrator that verifies and rewards sellers who reveal design flaws before a procurement

auction over ignoring the presence of potential flaws, using a simple price only auction, and renegotiating flaws after the allocation. Herweg and Schmidt (2018) assume that the cost of fixing a design flaw is higher after the assignment of the contract than before. This assumption drives the result that having an independent arbitrator find flaws before the allocation is preferable. However, should the buyer's choice of a procurement mechanism be unconstrained, it is not clear why it should be more expensive to fix design flaws after the allocation has been decided. An unexplored research question would be whether the common value aspect of the design flaws as well as the gains from allocative efficiency would be sufficient to generate a similar result like the one presented here.

Herweg and Schwarz (2018) analyse the problem of cost overruns in large multidimensional procurement projects when the buyer is constrained in the design of the procurement mechanism and an arbitrator is not available, and the procurer cannot commit not to renegotiate. The authors suggest a procurement mechanism that consists of two stages. In the initial stage, the procurer can extract rent from the firms through horizontal competition in a price only 'fixed design' auction. In the second stage, the procurer can achieve an efficient[3] design for the winning firm through Coasian bargaining. This mechanism can deliver both an efficient design and allocative efficiency, i.e. it selects the most productive firm and the most appropriate design for this firm. Herweg and Schwarz further argue that the procurer cannot benefit from allowing for multidimensional auctions in the first stage. However, this result is crucially dependent on the assumptions on the type of renegotiation in the second stage. Mainly, it is dependent on an implicit assumption that the procurer cannot effectively commit to distort their preferences over both phases of the procurement mechanism. Given the authors' assumption that the procurer cannot commit not to renegotiate in the second stage, this is a reasonable assumption, and the result stands. In chapter 2 and 3 of this thesis we introduce a related model of procurement with renegotiation, where the procurer has commitment power and gains from being able to use multidimensional auctions.

---

[3]Minimising the distance between social benefit and social cost

A seminal result of auction theory is that auctions strongly outperform negotiations. Bulow and Klemperer (1994) show that the addition of just one more bidder leads to a more substantial increase in payoffs for the auctioneer than the total gain made from moving from a simple to the optimal auction mechanism without the additional bidder. Herweg and Schmidt (2017) use their research on large procurement projects to show that in procurement, this is not as straightforward. Based on empirical findings of Bajari, R. McMillan, and Tadelis (2008), negotiations between the buyer and one seller are far more common in private procurement than auctions. To investigate why Herweg and Schmidt build a simple and compelling model of procurement with renegotiation. The authors show a trade-off between negotiations and auctions unexplored by Bulow and Klemperer (1994). Whereas the latter outperforms the former in the extraction of private rent, negotiations can improve on auctions by improving the communication between the two parties and thus facilitating less costly design improvements and improving overall efficiency. For complex procurement projects where scoring auctions are not feasible, negotiation between one seller and the buyer can often outperform price only auctions with renegotiation.

1.3.3 Long term relationships and non-contractible issues

We understand that in procurement settings, non-price elements may often play a crucial role. However, the buyer cannot always fully anticipate the importance of every element. Originally Hart, Shleifer, and Vishny (1997) and later Dalen, Moen, and Riis (2006) describe a process known as 'quality dumping', where firms exploit incomplete contracts by reducing an uncontracted quality aspect of the good in the production process to save costs. Hart, Shleifer, and Vishny (1997) use this to justify a trade-off between contracting out services and providing them in-house. Dalen, Moen, and Riis (2006) on the other hand, argue that endogenising the length of the contract can provide an effective incentive scheme against quality dumping in long-term dynamic procurement settings. In this way, "[...] contract-renewals

may be considered as an alternative to monetary rewards[...]" in long-term procurement settings (Dalen, Moen, and Riis, 2006).

In other cases, the buyer may not be able to fully contract each non-price element even when she does anticipate them. For instance, an element may not be verifiable by a third party, such as a court. Dellarocas, Dini, and Spagnolo (2006) argue that this can lead to two forms of opportunism: ex-ante and post-contracting opportunism. The former refers to situations where firms cannot credibly signal their desirable non-verifiable qualities leading to a mismatch. The latter refers to the 'quality dumping' described above. Dellarocas, Dini, and Spagnolo (2006) cite evidence that e-commerce marketplaces have adopted various kinds of reputation mechanisms to overcome these issues from opportunism. The authors provide an informal analysis of the mechanisms used by e-commerce providers at the time of writing. Calzolari and Spagnolo (2009) study the relationship between contract enforcement and competitive screening. Under the presence of non-contractibility issues the principal wants to restrict access to a subset of one or more of the agents and negotiates with these exclusively. In this way the principal trades off the gains from competition with gains from overcoming non-contractibility issues through 'loyalty'.

1.3.4 Repeated Interaction

Another interesting research avenue on the implication of long-term dynamics in procurement settings is the dynamics arising from repeated interaction between sellers and buyers rather than the dynamics arising within one interaction. Jofre-Bonet and Pesendorfer (2000) introduce a model of repeated procurement where some bidders are capacity constrained and where bidders can improve their technology through learning. Jofre-Bonet and Pesendorfer provide an empirical estimation of their model and show that there is indeed evidence for both the bidders' costs and the presence of capacity constraints affecting their bidding behaviour. This paper also gives an excellent introduction to the earlier literature on this

topic.

In much of the static procurement literature, it is implicitly assumed that different firms produce the same good, or that, at most, it might be the same good of different quality. In many environments, however, Lewis and Yildirim (2005) argue that even when goods are almost perfect substitutes, there may be costs to switch from one firm to another. For example, there may be a considerable cost for training and software, on top of the cost of new computers, to switch an economics department's computers from Microsoft to Apple and vice versa. In these environments, incumbent firms can price less competitively as the procurer is disincentivised to switch due to switching costs.

Lewis and Yildirim (2005) find that in procurement environments where the buyer is large enough to enjoy considerable market power, the result of the seller-buyer interaction may be very different. With commitment power, a large buyer may credibly threaten to switch from the incumbent seller to a new one and incur switching costs. Since the incumbent should realise that switching back will be made more expensive due to switching costs, he has a strong incentive to lower prices to maintain the relationship with the buyer. Lewis and Yildirim (2005) show in a simple model of switching cost that this effect may be strong enough to cause the buyer to achieve lower prices than she might in an environment without switching costs. Lewis and Yildirim (2006) summarise the findings from their joint research on switching costs in procurement. Procurers should generally prefer to stay with the incumbent. Except if by committing to switch unless their demands are met they can credibly force the incumbent to reduce prices - but as a trade-off will have to switch if he does not. This is work closely related to the wider literature on switching costs. Farrell and Klemperer (2007) give a good summary of the critical issues in, and contributions to this literature.

When buyers and sellers interact repeatedly, it also creates scope for the buyer to try and influence the sellers' bidders through releasing information in the bidding stages. For example, the question arises if it is optimal for the buyer to release all, some, or no infor-

mation about the firms' bids of the previous procurement round. Thomas (2010) addresses this question in a stylised two-round model of procurement with two identical contracts. She finds that the buyer usually achieves the best price by revealing as little information as possible, and ideally procures the two contracts simultaneously. However, if the contracts are offered sequentially, Thomas shows conditions on the strategic behaviour of the agents for which the buyer prefers to reveal all of the available information. This is possible as it can reduce the winner's incentive to make a worse offer in the second round believing the other seller to likely be a high-cost type.

In procurement practice, past performance in contracts is often rewarded for the allocation of new contracts. Intuitively, the buyer can learn about the sellers' types as well as trustworthiness over non-contractible issues and should reward those who perform well in the future to induce better performance during the contract. On the contrary, we may also argue that these rewards restrict entry and thus the buyer may suffer losses from reduced competition among the sellers. Butler et al. (2020) study this question in an experimental setting where the buyer sets varying rewards for past performance. They find that rewarding quality too much does indeed restrict entry. However, with appropriately balanced rewards for past performance both entry and quality provision can be improved for the buyer. However, in an empirical study of scoring rule auctions in Italian canteen services, Camboni, Valbonesi, et al. (2019) find that if the auction is rewarded to incumbents over new entrants, on average the buyer overpays for the good provision. This can be interpreted as evidence that in practice buyers overvalue the past performance of incumbents and are too risk-adverse in selecting a new, perhaps unknown entrant.

The study of repeated procurement auctions finally also allows for an analysis of the evolving dynamics of corruption. Chassang and Ortner (2019) develop a dynamic model of corruption in repeated procurement with minimum prices. In the context of procurement, this minimum price acts as a price-ceiling would in a regular auction. A standard result in auction theory is that in competitive environments, a minimum price should either not

affect the outcome of the procurement auction or make the outcome worse for the procurer. However, Chassang and Ortner (2019) show that under the presence of a cartel, a minimum price can improve the outcome for the procurer in expectation. This counterintuitive result is due to the constraints lowering the ability of the cartel to use price wars to punish defectors. Chassang and Ortner further support this argument with evidence from repeated procurements in Japanese procurement auctions. They find that within a prefecture previously shown to have a considerable presence of collusion by Ishii (2008), when a procurement auction had a minimum price, the impact of collusion is reduced, all else equal.

1.4 Dynamics for Key Contributions

In the previous subsection, we summarised a broad sample of important contributions to understanding how dynamics affect our predictions for the procurement environment. To put these into context, in the following section, we discuss how our understanding of the core principles of procurement should be affected by the papers discussed. How would the seminal papers of section 1.2 be affected if they take into account the lessons learnt from dynamic procurement? Furthermore, which, if any, of the results from the static literature are validated in the dynamic setting?

1.4.1 The Optimal Mechanism

The optimal static mechanism for multidimensional procurement described by Riordan and Sappington (1987a), McAfee and J. McMillan (1987), and Laffont and Tirole (1987) has been under much scrutiny in the dynamic procurement literature. Laffont and Tirole (1990) showed that even with no external impetus for renegotiation via some kind of shock, renegotiation dynamics considerably alter the outcome of a multidimensional auction. In this context, as we learnt above, the buyer-optimal mechanism distorts the non-price element away from the social optimum to limit the sellers' informational rent. Given full commit-

ment and no additional generation of private information after the initial stage, the optimal long-run contract thus also distorts these non-price elements such as quality or effort. Dewatripont (1989) showed that these contracts are not renegotiation-proof. In the context of procurement, this means that ex-post gains from trade can always be made if the buyer and seller agree to renegotiate the distorted buyer-optimal contract to socially optimal ones. Laffont and Tirole (1990) describe the optimal renegotiation-proof contract for the case of no commitment. They show that it implements the social optimum rather than the buyer optimum with full commitment.

Another question on the optimal mechanism is how it is affected if new private information is sequentially generated rather than available ex-ante. In the context of single-dimensional auctions, Eső and Szentes (2007) show that private information generated after the initial stage (second stage private information) does not create rent for the agents who hold it in the optimal dynamic mechanism. However, when private information is generated sequentially, and the environment is multidimensional, the second stage private information can affect the agents' rents (chapter 2). Thus for many procurement contexts, the seminal multidimensional optimal mechanism from Riordan and Sappington (1987a), McAfee and J. McMillan (1987), and Laffont and Tirole (1987) may not be the relevant guideline for a dynamic implementation. Finally, we saw above that Bajari, Tadelis, et al. (2001) argue that looking at the usually dynamic procurement environment through the lens of an asymmetric information problem is misleading. If their argument holds, and there is some evidence in support of it, looking at the procurement problem with mechanism design may be the wrong approach entirely.

In summary, therefore, from a perspective of mechanism design, whether a procurement setting is more accurately described as static or dynamic considerably changes both the environment and therefore the resulting optimal mechanism. Although some of the ideas from the static setting are validated, such as distorting the non-price element in multidimensional settings, we will see that many key results that rely on the optimal mechanism no longer hold

in the dynamic setting. For example, in chapter 2 of this thesis we see that the key result of Che (1993) about the optimality of scoring auctions no longer holds when the setting is dynamic.

1.4.2 The Optimal Scoring Auction

Given the popularity of scoring auctions in practice, we may wonder how their optimal design is affected by ex-post renegotiation. The crucial step to any extension of the seminal Che (1993) model to incorporate dynamics is the design of the renegotiation stage. In Herweg and Schwarz (2018), the procurer and seller use Coasian bargaining to agree on a price and quality in response to the cost shock. Herweg and Schwarz show that a multidimensional auction in the first stage as proposed by Che does not yield any benefit to the auctioneer. This result is a corollary of the more general result by Dewatripont (1989) that these types of auctions are not renegotiation proof. If the renegotiation stage for an extension of Che (1993) just maximises social surplus, then the Dewatripont result holds, and we cannot 'squeeze' the informational rent of agents in the first stage.

We propose another extension of Che (1993) to incorporate dynamics in chapter 3. Here the buyer commits to renegotiate only if the new price and quality still match the same score as the old ones at the original scoring rule. We argue that this commitment is a commitment to 'fairness', such that the buyer commits not to give the winning bidder advantages that the losing bidders did not have access to during the auction. In this environment, we show in chapter 3 that the scoring auction from Che (1993) is no longer able to implement the optimal mechanism with any scoring rule even with independent types. Bearing in mind the result from Branco (1997) that with correlated types, the scoring auction is also not optimal, these findings cast doubts on the performance of the scoring auction in practice. An extension of the Branco (1997) analysis to incorporate ex-post adaptations or renegotiation when types are correlated could make a valuable contribution to further the research on scoring auctions.

1.4.3 Adverse Selection

The Manelli and Vincent (1995) paper presents an insightful alternative angle to the analysis of multidimensional procurement as an adverse selection model which supports the conclusion that auctions may not be optimal. Dynamic models of adverse selection exist in other literatures (Guerrieri and Shimer, 2014; Hendel, Lizzeri, and Siniscalchi, 2005). Johnson (2013) provides a somewhat related model of adverse selection and dynamic moral hazard for procurement where the optimal dynamic contract can mitigate the inefficiency from adverse selection. These models show that considering dynamics can alter the predictions of adverse selection models considerably. However, in different contexts, for example in Bajari, Tadelis, et al. (2001), we find that auction mechanisms can become inefficient when there are considerable dynamics from renegotiation and add-on costs. This is supported by Doni (2006) who show that the auction mechanism further suffers from inefficiency in dynamic models as it leads to moral hazard issues. Whether the conclusion of the Manelli and Vincent (1995) adverse selection model on the inefficiency of auctions is retained in dynamic environments is thus an open question.

1.4.4 Corruption

The question of whether corruption in procurement is lessened with increased competition raised by Ganuza (2007) can be meaningfully extended by looking at the issue of corruption and renegotiation. Estache et al. (2009) argue that the presence of renegotiation increases the susceptibility of procurement to corruption. This argument is based on a common argument from auction theory that the presence of multiple bidders is in some way adverse to corruption.[4] Renegotiation, however, is a bilateral game and as such, lends itself more to corruption (Estache et al., 2009). Estache et al. (2009) find that anti-corruption policies are correlated with a significantly lower frequency of renegotiation. Estache et al. interpret this

---

[4]See for example Krishna (2009).

as evidence for a link between renegotiation and corruption. Guasch, Straub, and Laffont (2003) support this link and find empirical evidence for the theoretical model of Guasch, Laffont, and Straub (2006) that renegotiation is a factor that can lead to corruption. If the increase in corruption is due to the bilateral nature of renegotiation, this evidence supports the auction as a mechanism against corruption and goes against the arguments developed in Ganuza (2007).

Another avenue through which dynamics can impact the presence of corruption in procurement is through the repeated nature of many procurement environments. Factors such as reputation effects as discussed by Doni (2006) have an ambiguous effect. They could lead to increased corruption through increased rents from corruption in an environment with repeated interaction. Conversely, reputation punishments from corruption can also have a strong deterrent effect. Chassang and Ortner (2019) discuss a model of repeated procurement with cartels. Here the auctioneer has the option to set a minimum price, which is equivalent to a maximum price in a regular auction and would not usually be beneficial in an environment without corruption. However, Chassang and Ortner show that this mechanism constrains cartels in their scope for punishing defectors. In this way, depending on the extent of corruption, a minimum price can make the auctioneer better off. The authors support this theory by drawing on evidence from Japanese procurement where they show that a minimum price is associated with reduced levels of corruption in the data.

1.5   Conclusion

The procurement environment has received considerable attention in the economic literature over recent years. This paper introduced a summary of a selection of these contributions. We summarised seminal works that formed our understanding of procurement today and important new ideas originating from the real-life observation that procurement is hardly ever a static game. Finally, we considered how these new works may influence the way we

should interpret the results of the seminal works.

Central to the theory of procurement are developments in mechanism design. From the early procurement literature centred around the work of Laffont and Tirole in the 1980s, the discovery of the optimal mechanism for multidimensional environments by McAfee and J. McMillan (1987), Riordan and Sappington (1987a), and Laffont and Tirole (1987), and the implementation of it through scoring auctions, mechanism design lies at the heart of many of the key contributions to our understanding of procurement. Perhaps unsurprisingly then, the recent development of considering dynamic issues to procurement coincides with a surge in our comprehension of optimal dynamic mechanism design. Works such as Arve and Martimort (2016), or the chapters 2 and 3 of this thesis build on and contribute to this field.

As discussed in this review, the branch of the literature centred around the work from Bajari and Tadelis provides a critical perspective on the strong link between procurement and mechanism design. By reducing this environment to a game of private information, important ideas may be lost. In particular, Bajari and Tadelis argue that the mechanism design approach to procurement has failed to properly take into account many of the relevant considerations for firms arising from the frequency that procurement contracts are renegotiated. The literature on dynamic procurement in part addresses these concerns, but a disconnect still persists. Renegotiation in practice is often a costly inefficiency that necessarily arises due to incomplete contracting. However, in many theoretical works, this is not reflected well. In some cases, the renegotiation is instead considered the opposite - a costless adaptation tool to exploit gains from trade arising through the inherent dynamics of the model. A convincing model of procurement that addresses these concerns is perhaps one of the biggest research gaps left to explore.

Chapter Two

RENEGOTIATION IN PROCUREMENT - DYNAMIC MULTIDIMENSIONAL AUCTION DESIGN

2.1 Introduction

Procurement today makes up a large share of developed nations' GDP. Public government procurement alone accounted for 12% of GDP across OECD nations over the last decade (OECD, 2015). The benefits of purchasing through procurement are well documented, and evidence from the private and public sector is readily available.

The design of the procurement process is of critical importance to both governments and private institutions seeking to undertake a procurement project. As such, this process has been studied extensively in economic theory. Several aspects make procurement a particularly interesting topic to study from the standpoint of theory. For one, goods or services to be procured tend to have multiple dimensions of interest in addition to merely the price. Depending on the application, the timing of delivery, the reliability of the product, or even

the design may be as important to the procurer as the price. Even procurement contracts for seemingly simple goods such as latex gloves for hygiene and protection can involve multiple quality considerations such as the thickness of the gloves and their durability. Furthermore, as the procurement process can be lengthy and expensive, procurers usually prefer to set out long term contracts which introduce dynamic considerations. Contracts have to be set to allow for various contingencies and the contract and its terms may be changed and renegotiated.

In this paper, one particularly popular procurement process is analysed: the scoring auction. A scoring auction is a reverse auction format in which the sellers' bids are ranked by scores determined by a scoring rule that rewards price and non-price elements of the procurement project. The procurer chooses the scoring rule to reflect their preferences over the price and non-price aspects. Scoring auctions provide a decision rule to compare complex and possibly widely different contract offers and are commonly applied in practice. In economics these auctions have been studied among others by Che (1993), Branco (1997), Chao and Wilson (2002), and Asker and Cantillon (2008). This paper adds to the existing literature by analysing how the design of the optimal scoring auction changes when allowing for the possibility of post-allocation renegotiation.

The prevalence of renegotiation in procurement is partly a consequence of two of the distinguishing characteristics of procurement processes described above. As procurement contracts are usually long term, unforeseen contingencies such as cost shocks, changes in preferences over non-price aspects, or technology developments, may occur during the term of the supply contract. Furthermore, the multidimensionality of the contract allows for the procurer and the supplier to renegotiate the terms of the different dimensions in response to these shocks. There are thus two motivations for renegotiation in multidimensional procurement, one negative and one positive. In long term contracts, unexpected shocks that were not contracted on inevitably arise. However, given a shock to the cost of producing a non-price element, *both* the seller and the buyer can usually benefit from reagreeing the element

In this way gains from trade can be realised that are not available in a single-dimensional environment.

In this paper, the model environment of Che (1993) on scoring auctions and procurement is extended to include uncertainty for the seller about the cost of procuring a good. This extension allows us to address several research questions. First and foremost, we want to understand whether scoring auctions can implement the optimal mechanism when the contracts are renegotiated in response to these cost shocks. Here we describe the properties of the buyer-optimal mechanism and show that scoring auctions with renegotiation (SAWR) are not able to implement it.

The optimal mechanism problem in this setting has several interesting features. The problem is multidimensional so that a social choice function specifies quality policies in addition to transfers and an allocation rule for each agent. Furthermore, the problem is dynamic and includes two stages, an allocation, and a renegotiation stage, with sequential generation of private information. Finally, to match the feature of a procurer renegotiating with only one agent, the problem restricts participation in the renegotiation phase to only the agent winning the allocation decision. Laffont and Tirole (1987), McAfee and J. McMillan (1987), and Riordan and Sappington (1987a) extend the Myerson (1981) framework to a multidimensional setting. Unlike these papers, we consider a multidimensional and dynamic auction setting. This paper thus relates to Eső and Szentes (2007), Pavan, Segal, and Toikka (2014) and the literature on dynamic mechanism design. Eső and Szentes describe the optimal mechanism for a two-stage single-dimensional (price only) auction with all bidders taking part in both stages, whereas here the setting is of a multidimensional reverse auction with all but one bidders being excluded from the second stage (renegotiation). Pavan, Segal, and Toikka (2014) describe necessary conditions for incentive compatibility and the use of the envelope formula in general dynamic mechanism design problems. Unlike these seminal contributions to dynamic mechanism design, our analysis maintains a static choice rule over the allocation of the good to an agent and generates dynamics through the multidimension-

ality of the agents' preferences. In the sense that only one bidder remains in the second stage and the project is multidimensional, this work is thus more closely related to Riordan and Sappington (1987b). However, where Riordan and Sappington are concerned with how the choice of optimal operational mode is driven by the correlation of the private information in the two stages of a principal-agent model. Whereas here the focus lies on the optimal design of the allocation mechanism to one of several agents.

Some notable research on procurement and renegotiation or procurement with cost uncertainty was done by Tirole (1986), Laffont and Tirole (1990), and more recently Arve and Martimort (2016) as well as Herweg and Schmidt (2018). Tirole (1986) focus on the issue of adverse selection given an incomplete contract whereas we study complete contracts. Laffont and Tirole (1990) study the implications of lack of commitment not to renegotiate whereas we study renegotiation under full commitment. Arve and Martimort (2016) study how the buyer can exploit the risk aversion of agents for the design of the optimal contract. We assume risk neutrality throughout and study the impact of sequential private information on the rent distortion through an optimal mechanism. Finally, Herweg and Schmidt (2018) restrict attention to particularly large and complex procurement projects whereas we make no such assumption. Other related literature to our work is done by Eső and Szentes (2007) on dynamic mechanisms and Pai and Vohra (2013) as well as Quint and Hendricks (2018) on dynamic auctions. We extend the work of Eső and Szentes (2007) to allow for multidimensional type spaces and focus in particular on dynamic procurement auctions. Pai and Vohra (2013) focus on dynamic auctions with bidders arriving over time and the multidimensionality of the auction stemming from changing arrival and departure times. In our case all bidders arrive at the same time and the multidimensionality is inherent to the good and not just the bidder. In Quint and Hendricks (2018), the auction is dynamic due to entry costs but otherwise single-dimensional. In our analysis there are no entry costs and the renegotiation dynamics arise due to the multidimensionality of the good. It is to the best of our knowledge fully unexplored in the literature how the design of the scoring auction is affected by the

possibility of renegotiation. The main contributions of this paper are thus twofold. First, we show that when allowing for the possibility of renegotiation, the result of Che (1993) breaks down and scoring auctions are no longer able to implement the optimal mechanism. Second, we contribute to the literature on dynamic mechanism design by outlining the solution to the multidimensional dynamic mechanism design programme detailed above. Here we show that unlike in Eső and Szentes (2007) the second stage private information does affect the informational rent an agent can extract.

This paper is organised as follows. Section 2.2 outlines the model. Section 2.3 sets up the mechanism design problem and describes its solution. Section 2.4 introduces a scoring auction with renegotiation (SAWR) implementation game and shows that it cannot implement the social choice function of the optimal mechanism design problem from section 2.3. The appendices contain proofs and derivations for various theorems and lemmas introduced throughout the analysis.

2.2   The Model

N agents indexed $i = \{1, ..., N\}$ are competing to sell an indivisible good to the principal. Agent i has cost-type $\theta_i \in \Theta$ drawn (i.i.d.) from a continuous distribution with density $f(\theta)$ and cumulative density $F(\theta)$. Let $[\underline{\theta}, \overline{\theta}]$ be the support of this distribution where $\underline{\theta}$ is the lowest θ with a positive density and $\overline{\theta}$ is the highest. If agent i supplies the good at quality q for a price p, the principal obtains utility $U(p, q) = V(q) - p$. The agent makes profits π from (p, q). Here $\pi = p - c(q, \theta_i, \varepsilon)$ is the difference between the transfers p from the principal to the agent and the cost $c(q, \theta_i, \varepsilon)$ to the agent of producing the good at quality q. The cost depends on the agent's cost parameter θ_i and a parameter ε. This ε denotes the realisation of a shock drawn from a continuous distribution with probability density $g(\varepsilon)$ and cumulative density $G(\varepsilon)$ on the support $[\underline{\varepsilon}, \overline{\varepsilon}]$. The realisation ε is drawn independently of the agents' realisations of types $\boldsymbol{\theta} = (\theta_i, \boldsymbol{\theta}_{-i})$ and is unobserved by the principal. The timing

of events is as follows:

1. The principal commits to a mechanism.

2. The agents privately learn their cost type θ.

3. The agents choose a strategy denoting messages at each node of the game tree of the mechanism.

4. The mechanism may allocate the procurement of the good to an agent given the agents' messages.

5. If and only if an agent has been allocated the procurement contract, he privately realises the shock ε and thus knows his true cost parameter θ, ε. He may choose any additional message after the shock in response to this new private information.

6. The agent from step 5 produces the good for the principal given the outcome function and the agents' messages before and after the shock.

The principal's preferences over quality $V(q)$ satisfies $V'(q) > 0$ and $V''(q) < 0$ ($\forall q > 0$) and the Inada conditions $\lim_{q \to 0} V'(q) = \infty$ and $\lim_{q \to \infty} V'(q) = 0$ to guarantee an interior solution. The cost functions are increasing in all arguments everywhere and are convex in quality: $c_1 > 0$, $c_2 > 0$, $c_3 > 0$, $c_{11} \geq 0$. The cross derivatives additionally satisfy: $c_{12} \geq 0$ and $c_{13} \geq 0$, i.e. marginal cost is increasing in the cost-type as well as the shock.[1] We additionally assume all agents and the principal to be risk neutral. This assumption simplifies the division of the two stages of the mechanism before and after the shock. Since agents may face an ex-post loss, without risk neutrality the analysis is considerably more complicated. Intuitively, we can think of risk-neutral sellers and buyers as profit maximising agents for which the size of the procurement contract is small relative to the overall firm.

---

[1]Functions that satisfy these conditions and the model assumptions are for example $(\theta + \varepsilon)^k q$ for $k \geq 1$ for the cost functions and q^l for $l \in (0, 1)$ for the utility function of the principal.

If we consider a smaller firm then the assumption of risk neutrality might not be tenable since with a large procurement contract, making an ex-post loss might create a possibility of bankruptcy. The following additional assumptions simplify the analysis considerably:

Assumption 1. Monotone Hazard Rates *Both $F(\theta)$ and $G(\varepsilon)$ satisfy the increasing hazard rate property, i.e. $F(\theta)/f(\theta)$ is nondecreasing in θ and $G(\varepsilon)/g(\varepsilon)$ is nondecreasing in ε.*

Assumption 2. Dynamic Regularity $c_1 + \frac{F}{f}(c_{12} - \frac{G}{g}c_{123})$ *is nondecreasing in θ, ε, and q.*

This assumption is to ensure monotonicity in the optimal quality policies q for both θ and ε and to guarantee concavity of the optimisation programme in q. It serves the same purpose as assumption 1 in Che (1993) and works in conjunction with assumption 1 to rule out some very steep cost functions such as $(\theta+\varepsilon)^k q$ with large k. This is a strong assumption that rules out many reasonable cost functions. It, or a functionally equivalent assumption, is necessary for the problem to be monotonic and for us to describe an analytical solution.

Assumption 3. *There are gains from trade to be made even with a firm of the highest cost type $\overline{\theta}$ and the highest possible cost shock $\overline{\varepsilon}$, i.e. for some q, $V(q) > c(q,\overline{\theta},\overline{\varepsilon})$.*

2.3   The Optimal Mechanism

The principal's objective is to maximise her expected utility by designing a mechanism that defines to every agent for every realisation of agents' strategies a probability of being allocated the contract x_i, transfers to the agent from the principal t_i, and a quality q_i to be implemented by the agent if they win the allocation decision. Of interest are only mechanisms in which the allocation choice is made before the realisation of the shock to the cost. In other words, the t_i and q_i may be contingent on the agents' strategies *after* the realisation of the shock ε, but the x_i cannot. By an extension of the revelation principle, the strategies available to the agents can be simplified, and attention can be restricted to mechanisms in which the agents

simultaneously, directly, and truthfully reveal their types and the winning agent reveals the cost shock directly and truthfully:

Lemma 2.3.1. A Revelation Principle *Any perfect Bayesian equilibrium outcome of any mechanism of the model setup described in section 2.2 can be implemented by a direct mechanism in which the N initial agents simultaneously and truthfully reveal their type θ to the principal and the one remaining agent truthfully reveals ε to the principal.*

See appendix A.1 for a full proof of this extension of the revelation principle and a full definition of a direct mechanism in this environment. There are thus two stages to the direct mechanism. The first stage comprises of an initial transfer and a decision rule for the allocation of the second stage. In the subsequent stage, there may be further transfers and the qualities are chosen in response to the revelation of the shock.

2.3.1 The Principal's Optimisation Programme

The principal's problem is then to choose $\{x_i(\cdot), t_i^1(\cdot), t_i^2(\cdot, \cdot), q_i(\cdot, \cdot)\}$ for all i, $\boldsymbol{\theta}$, and ε to maximise her expected utility subject to incentive and participation constraints. Here t_i^1 and t_i^2 denote the first and second stage transfers, respectively. The agent i's total transfers from the mechanism are denoted $t_i(\boldsymbol{\theta}, \varepsilon) = t_i^1(\boldsymbol{\theta}) + t_i^2(\boldsymbol{\theta}, \varepsilon)x_i(\boldsymbol{\theta})$. Also let: $\boldsymbol{t}(\boldsymbol{\theta}, \varepsilon) = (t_1(\boldsymbol{\theta}, \varepsilon), \cdots, t_N(\boldsymbol{\theta}, \varepsilon))$, $\boldsymbol{x}(\boldsymbol{\theta}) = (x_1(\boldsymbol{\theta}), \cdots, x_N(\boldsymbol{\theta}))$, and $\boldsymbol{q}(\boldsymbol{\theta}, \varepsilon) = (q_1(\boldsymbol{\theta}, \varepsilon), \cdots, q_N(\boldsymbol{\theta}, \varepsilon))$ for all $\boldsymbol{\theta}, \varepsilon$. We initially assume that the principal chooses not to reveal to the winning agent any information about the other agents' types. This restriction will turn out to not be restrictive. The principal's problem is thus to solve:

$$\max_{\{\boldsymbol{x}(\boldsymbol{\theta}), \boldsymbol{t}(\boldsymbol{\theta}, \varepsilon), \boldsymbol{q}(\boldsymbol{\theta}, \varepsilon)\}} \sum_{i=1}^{N} E_{\varepsilon, \boldsymbol{\theta}} \Big[V(q_i(\boldsymbol{\theta}, \varepsilon))x_i(\boldsymbol{\theta}) - t_i(\boldsymbol{\theta}, \varepsilon) \Big], \tag{2.1}$$

subject to the constraints of first stage participation:

$$E_{\varepsilon, \boldsymbol{\theta}_{-i}} \Big[t_i(\boldsymbol{\theta}, \varepsilon) - c(q_i(\boldsymbol{\theta}, \varepsilon), \theta_i, \varepsilon)x_i(\boldsymbol{\theta}) \Big] \geq 0 \tag{2.2}$$

34

$(\forall i, \theta_i)$, first stage truthtelling:

$$\theta_i \in \operatorname*{argmax}_{\widetilde{\theta}_i} E_{\varepsilon, \boldsymbol{\theta}_{\text{-}i}} \left[v_i(\widetilde{\theta}_i, \theta_i, \varepsilon) x_i(\widetilde{\theta}_i, \boldsymbol{\theta}_{\text{-}i}) + t_i^1(\widetilde{\theta}_i, \boldsymbol{\theta}_{\text{-}i}) \right] \tag{2.3}$$

$(\forall i, \theta_i)$, where:

$$v_i(\widetilde{\theta}_i, \theta_i, \varepsilon) = \max_{\widetilde{\varepsilon}} E_{\boldsymbol{\theta}_{\text{-}i}} \left[t_i^2(\widetilde{\theta}_i, \boldsymbol{\theta}_{\text{-}i}, \widetilde{\varepsilon}) - c(q_i(\widetilde{\theta}_i, \boldsymbol{\theta}_{\text{-}i}, \widetilde{\varepsilon}), \theta_i, \varepsilon) \right] \tag{2.4}$$

$(\forall i, \theta_i, \widetilde{\theta}_i)$, and second stage truthtelling:

$$\varepsilon \in \operatorname*{argmax}_{\widetilde{\varepsilon}} E_{\boldsymbol{\theta}_{\text{-}i}} \left[t_i^2(\theta_i, \boldsymbol{\theta}_{\text{-}i}, \widetilde{\varepsilon}) - c(q_i(\theta_i, \boldsymbol{\theta}_{\text{-}i}, \widetilde{\varepsilon}), \theta_i, \varepsilon) \right] \tag{2.5}$$

$(\forall i, \theta_i, \varepsilon)$. In this formulation of the incentive constraints, the possibility of double deviations is taken into account. I.e. if the constraints (2.3) and (2.5) are satisfied it cannot be profitable (in expectation) for an agent to misreport θ in the first stage, planning to misreport ε in the second stage. As such, the second stage truthtelling constraint only needs to be satisfied given truthtelling in the first stage. $v_i(\widetilde{\theta}_i, \theta_i, \varepsilon)$ denotes the expectation over the value of reaching the second stage and realising shock ε given a type θ_i and first stage report $\widetilde{\theta}_i$.

It's possible here to draw comparisons to the model environment of other analyses to better place this paper in the literature. First, if the second stage is ignored, the model setup here is identical to Che (1993). Second, more interestingly, if $N = 1$, the model setup is similar to the setup of Riordan and Sappington (1987b), but with a unit quantity and without the principal being given the option to acquire the ability to produce the good themselves by buying up the agent. This last difference is due to the focus of Riordan and Sappington (1987b) on vertical integration (organisational mode), which is not within the scope of the analysis here. Finally, there are some parallels to the setup of Eső and Szentes (2007). Aside from being single rather than multidimensional, as their focus is not on renegotiation, Eső and Szentes (2007) allow all N agents to partake in the second stage of the mechanism. The second stage of their model thus enables the principal to realise efficiency gains by potentially choosing a different (ex-post more efficient) agent rather than

the most efficient one in expectation. In contrast, in the multidimensional setting of this analysis, the second stage allows the principal to realise efficiency gains by choosing a different price and quality combination in response to the shock. This type of second stage is only made possible by the multidimensionality of the environment.

2.3.2 The Optimal Direct Mechanism

To proceed, we provide a step-by-step solution to the direct mechanism optimisation programme. First, we can simplify the principal's optimisation programme with the following lemma:

Lemma 2.3.2. *The optimal quality policies $\boldsymbol{q}$ which form part of the solution to the principal's problem from section 2.3.1 do not depend on the realisations of θ of the $N-1$ agents who do not win the allocation.*

See appendix A.2 for a proof of this lemma, where we use Jensen's inequality to show that when an agent's quality policies depend on the realisation of other agents' types, the principal can improve the optimisation programme, leading to a contradiction. This proof is adapted from Laffont and Tirole (1987) who also give an intuitive justification of the result. We can thus write $q_i(\theta_i, \boldsymbol{\theta}_{\text{-}i}, \varepsilon)$ as $q_i(\theta_i, \varepsilon)$ without loss of generality.

Corollary 2.3.2.1. *The second stage transfers $t_i^2(\theta_i, \boldsymbol{\theta}_{\text{-}i}, \varepsilon)$ can be restricted to not depend on the other agents' types without loss of generality.*

This result is immediate from Lemma 2.3.2 and the risk neutrality of the agents. We can thus write $t_i^2(\theta_i, \boldsymbol{\theta}_{\text{-}i}, \varepsilon)$ as $t_i^2(\theta_i, \varepsilon)$ without loss of generality. Since all second stage variables do not depend on the other agents' types, our initial assumption that the principal does not disclose any information about these to the winning agent was not restrictive. We can now obtain an expression of the agent's value of participating in the second stage for a given shock ε.

Lemma 2.3.3. *By the envelope theorem, we can equivalently write the agent i's expected value of the second stage subject to second stage truthtelling from equation (2.5) as:*

$$v_i(\widetilde{\theta}_i, \theta_i, \varepsilon) = \int_{\varepsilon}^{\overline{\varepsilon}} c_3(q_i(\widetilde{\theta}_i, s), \theta_i, s)ds + v_i(\widetilde{\theta}_i, \theta_i, \overline{\varepsilon}) \tag{2.6}$$

($\forall i, \widetilde{\theta}_i, \theta_i, \varepsilon$ such that $\widetilde{\theta}_i = \theta_i$).

See appendix A.3 for a full derivation of this result, where we assume the optimal quality policies to be non-increasing in ε, which will later be verified to be true. Using the envelope theorem in the same way as for lemma 2.3.3, we can use equation (2.6) to find an expression for the expected value of participating in the mechanism for an agent i of type θ_i.

Lemma 2.3.4. *By the envelope theorem, given the first and second stage truthtelling constraints from equations (2.3) and (2.5) as well as the participation constraint (2.2), the value $W_i(\theta_i)$ for an agent i of type θ_i to participate in the mechanism can be written as:*

$$W_i(\theta_i) = E_{\varepsilon, \boldsymbol{\theta}_{-i}}\left[\int_{\theta_i}^{\overline{\theta}}\left(c_2(q_i(t, \overline{\varepsilon}), t, \overline{\varepsilon}) - \int_{\varepsilon}^{\overline{\varepsilon}} c_{32}(q_i(t, s), t, s)ds\right)x_i(t, \boldsymbol{\theta}_{-i})dt\right] \tag{2.7}$$

($\forall i, \theta_i$).

See appendix A.4 for the proof of lemma 2.3.4. Note that we can also express the value $W_i(\theta_i)$ of participating in the mechanism in truthtelling equilibrium as:

$$W_i(\theta_i) = E_{\varepsilon, \boldsymbol{\theta}_{-i}}\left[t_i(\theta_i, \boldsymbol{\theta}_{-i}, \varepsilon)\right] - E_{\varepsilon, \boldsymbol{\theta}_{-i}}\left[c(q_i(\theta_i, \varepsilon), \theta_i, \varepsilon)x_i(\theta_i, \boldsymbol{\theta}_{-i})\right] \tag{2.8}$$

($\forall i, \theta_i$). For a given $E_{\boldsymbol{\theta}_{-i}}[x_i(\theta_i, \boldsymbol{\theta}_{-i})]$ ex-ante symmetry allows us to decompose the optimisation problem from equation (2.1) into N problems of the type:

$$\max_{t_i(\boldsymbol{\theta}, \varepsilon), q_i(\theta_i, \varepsilon)} E_{\varepsilon, \theta_i, \boldsymbol{\theta}_{-i}}\left[V(q_i(\theta_i, \varepsilon))x_i(\theta_i, \boldsymbol{\theta}_{-i}) - t_i(\theta_i, \boldsymbol{\theta}_{-i}, \varepsilon)\right], \tag{2.9}$$

subject to (2.2), (2.3), and (2.5) ($\forall i$). We can now substitute for $E_{\varepsilon, \boldsymbol{\theta}_{-i}}\left[t_i(\theta_i, \boldsymbol{\theta}_{-i}, \varepsilon)\right]$ from equation (2.8) to get:

$$\max_{q_i(\theta_i, \varepsilon)} E_{\varepsilon, \theta_i}\left[\left(V(q_i(\theta_i, \varepsilon)) - c(q_i(\theta_i, \varepsilon), \theta_i, \varepsilon)\right)E_{\boldsymbol{\theta}_{-i}}\left[x_i(\theta_i, \boldsymbol{\theta}_{-i})\right] - W_i(\theta_i)\right] \tag{2.10}$$

37

($\forall i$). Finally, substituting for $W_i(\theta_i)$ from (2.7) in Lemma 2.3.4 and thus enforcing constraints (2.2), (2.3), and (2.5) we get:

$$\max_{q_i(\theta_i,\varepsilon)} E_{\varepsilon,\theta_i}\left[\left(V(q_i(\theta_i,\varepsilon)) - c(q_i(\theta_i,\varepsilon),\theta_i,\varepsilon)\right)E_{\boldsymbol{\theta}_{-i}}\left[x_i(\theta_i,\boldsymbol{\theta}_{-i})\right]\right.$$
$$\left. - \int_{\theta_i}^{\overline{\theta}}\left(c_2(q_i(t,\overline{\varepsilon}),t,\overline{\varepsilon}) - \int_{\underline{\varepsilon}}^{\overline{\varepsilon}} c_{32}(q_i(t,s),t,s)ds\right)E_{\boldsymbol{\theta}_{-i}}\left[x_i(t,\boldsymbol{\theta}_{-i})\right]dt\right] \qquad (2.11)$$

($\forall i$) as an equivalent expression of the principal's optimisation programme. Writing out the expectation and using Fubini and Tonelli's theorems for changing the order of integration, we can integrate by parts twice to simplify this messy expression. First[2], to get:

$$\max_{q_i(\theta_i,\varepsilon)} \int_{\underline{\varepsilon}}^{\overline{\varepsilon}}\int_{\underline{\theta}}^{\overline{\theta}}\left(\left(V(q_i(\theta_i,\varepsilon)) - c(q_i(\theta_i,\varepsilon),\theta_i,\varepsilon)\right)E_{\boldsymbol{\theta}_{-i}}\left[x_i(\theta_i,\boldsymbol{\theta}_{-i})\right]g(\varepsilon)\right.$$
$$\left. - \int_{\theta_i}^{\overline{\theta}}\left(c_2(q_i(t,\overline{\varepsilon}),t,\overline{\varepsilon})g(\varepsilon) - c_{32}(q_i(t,\varepsilon),t,\varepsilon)G(\varepsilon)\right)E_{\boldsymbol{\theta}_{-i}}\left[x_i(t,\boldsymbol{\theta}_{-i})\right]dt\right)f(\theta_i)d\theta_i d\varepsilon. \qquad (2.12)$$

And second[3], to get:

$$\max_{q_i(\theta_i,\varepsilon)} \int_{\underline{\varepsilon}}^{\overline{\varepsilon}}\int_{\underline{\theta}}^{\overline{\theta}}\left(\left(V(q_i(\theta_i,\varepsilon)) - c(q_i(\theta_i,\varepsilon),\theta_i,\varepsilon)\right)g(\varepsilon)f(\theta_i)\right.$$
$$\left. - \left(c_2(q_i(\theta_i,\overline{\varepsilon}),\theta_i,\overline{\varepsilon})g(\varepsilon) - c_{32}(q_i(\theta_i,\varepsilon),\theta_i,\varepsilon)G(\varepsilon)\right)F(\theta_i)\right)E_{\boldsymbol{\theta}_{-i}}\left[x_i(\theta_i,\boldsymbol{\theta}_{-i})\right]d\theta_i d\varepsilon. \qquad (2.13)$$

As everything is additively separable, maximising this double integral is equivalent to pointwise maximising its integrand. This quality also allows us to treat $E_{\boldsymbol{\theta}_{-i}}\left[x_i(\theta_i,\boldsymbol{\theta}_{-i})\right]$ as given for the time being since, for any optimal decision rule $x_i(\theta_i,\boldsymbol{\theta}_{-i})$, it acts as a scalar on the integrand. Therefore:

$$\max_{q_i(\theta_i,\varepsilon)} \left(V(q_i(\theta_i,\varepsilon)) - c(q_i(\theta_i,\varepsilon),\theta_i,\varepsilon)\right)g(\varepsilon)f(\theta_i)$$
$$- \left(c_2(q_i(\theta_i,\overline{\varepsilon}),\theta_i,\overline{\varepsilon})g(\varepsilon) - c_{32}(q_i(\theta_i,\varepsilon),\theta_i,\varepsilon)G(\varepsilon)\right)F(\theta_i) \qquad (2.14)$$

---

[2]When $u = \int_{\varepsilon}^{\overline{\varepsilon}} c_{32}(q_i(t,s),t,s)ds$ and $v' = g(\varepsilon)$, then $\int_{\underline{\varepsilon}}^{\overline{\varepsilon}} uv'd\varepsilon = \left[uv\right]_{\underline{\varepsilon}}^{\overline{\varepsilon}} - \int_{\underline{\varepsilon}}^{\overline{\varepsilon}} u'vd\varepsilon$. Thus $\int_{\underline{\varepsilon}}^{\overline{\varepsilon}}\int_{\varepsilon}^{\overline{\varepsilon}} c_{32}(q_i(t,s),t,s)dsg(\varepsilon)d\varepsilon = \left[\int_{\varepsilon}^{\overline{\varepsilon}} c_{32}(q_i(t,s),t,s)dsG(\varepsilon)\right]_{\underline{\varepsilon}}^{\overline{\varepsilon}} + \int_{\underline{\varepsilon}}^{\overline{\varepsilon}} c_{32}(q_i(t,\varepsilon),t,\varepsilon)G(\varepsilon)d\varepsilon$. This is the same as $\int_{\underline{\varepsilon}}^{\overline{\varepsilon}} c_{32}(q_i(t,\varepsilon),t,\varepsilon)G(\varepsilon)d\varepsilon$, since the first part of the expression is 0.

[3]Let $u = \int_{\theta_i}^{\overline{\theta}}\left(c_2(q_i(t,\overline{\varepsilon}),t,\overline{\varepsilon})g(\varepsilon) - c_{32}(q_i(t,\varepsilon),t,\varepsilon)G(\varepsilon)\right)E_{\boldsymbol{\theta}_{-i}}[x_i(t,\boldsymbol{\theta}_{-i})]dt$ and $v' = f(\theta_i)$. Then follow the same steps as before.

is an equivalent representation of the principal's problem for any given optimal $E_{\boldsymbol{\theta}_{\text{-}i}}\left[x_i(\theta_i, \boldsymbol{\theta}_{\text{-}i})\right]$. Finally, from ex-ante symmetry, the subscript i can be dropped and the optimisation programme may be rewritten as:

$$\max_{q(\theta,\varepsilon)}\Big(V(q(\theta,\varepsilon)) - c(q(\theta,\varepsilon),\theta,\varepsilon)\Big)g(\varepsilon)f(\theta)$$

$$-\Big(c_2(q(\theta,\overline{\varepsilon}),\theta,\overline{\varepsilon})g(\varepsilon) - c_{32}(q(\theta,\varepsilon),\theta,\varepsilon)G(\varepsilon)\Big)F(\theta). \tag{2.15}$$

Solving the principal's problem for the optimal quality policies $q(\theta,\varepsilon)$ thus involves solving (2.15) with respect to q. The first-order conditions of the optimisation problem give necessary and sufficient[4] conditions for the optimal qualities:

$$V'(q(\theta,\varepsilon) = c_1(q(\theta,\varepsilon),\theta,\varepsilon)g(\varepsilon)f(\theta) - c_{21}(q(\theta,\overline{\varepsilon}),\theta,\overline{\varepsilon})g(\varepsilon)F(\theta)$$

$$- c_{321}(q(\theta,\varepsilon),\theta,\varepsilon)G(\varepsilon)F(\theta) \tag{2.16}$$

$(\forall \theta, \varepsilon)$. From assumption 2, $q(\theta_i,\varepsilon)$ is a nonincreasing function of θ_i and ε. Given the maximand from equation (2.15) is concave in q and q is nonincreasing in θ, the maximand is decreasing in θ. Thus a lower θ is worth strictly more selecting, and the optimal $x_i(\boldsymbol{\theta})$ are:

$$x_i(\theta_i, \boldsymbol{\theta}_{\text{-}i}) = 1 \text{ if } \theta_i < \min \boldsymbol{\theta}_{\text{-}i} \tag{2.17}$$

$$\text{and } x_i(\theta_i, \boldsymbol{\theta}_{\text{-}i}) = 0 \text{ if } \theta_i > \min \boldsymbol{\theta}_{\text{-}i} \tag{2.18}$$

$(\forall i, \boldsymbol{\theta})$. Substituting equation (A.29) into equation (2.8) we can get an expression for the expected equilibrium transfers of agent i:

$$E_{\varepsilon,\boldsymbol{\theta}_{\text{-}i}}\Big[t_i(\theta_i, \boldsymbol{\theta}_{\text{-}i}, \varepsilon)\Big] = E_{\varepsilon,\boldsymbol{\theta}_{\text{-}i}}\bigg[\int_{\theta_i}^{\overline{\theta}}\Big(c_2(q_i(t,\overline{\varepsilon}),t,\overline{\varepsilon}) - \int_{\varepsilon}^{\overline{\varepsilon}} c_{32}(q_i(t,s),t,s)ds\Big)x_i(t,\boldsymbol{\theta}_{\text{-}i})dt$$

$$+ c(q_i(\theta_i,\varepsilon),\theta_i,\varepsilon)x_i(\theta_i, \boldsymbol{\theta}_{\text{-}i})\bigg] \tag{2.19}$$

---

[4]Sufficient, since the maximand is concave in q (from assumption 2) and the assumptions on the principal's preferences over quality guarantee an interior solution.

$(\forall i, \theta_i)$. Given equations (2.17) and (2.18), we can rewrite this equation as:

$$E_{\varepsilon, \boldsymbol{\theta}_{-i}}\left[t_i(\theta_i, \boldsymbol{\theta}_{-i}, \varepsilon)\right] = E_\varepsilon\left[\int_{\theta_i}^{\overline{\theta}}\left(c_2(q_i(t, \overline{\varepsilon}), t, \overline{\varepsilon}) - \int_\varepsilon^{\overline{\varepsilon}} c_{32}(q_i(t, s), t, s)ds\right)X_i(t)dt\right.$$
$$\left. + c(q_i(\theta_i, \varepsilon), \theta_i, \varepsilon)X_i(\theta_i)\right] \tag{2.20}$$

$(\forall i, \theta_i)$, where $X_i(\theta_i)$ is the probability $Pr(\theta_i < \min \boldsymbol{\theta}_{-i})$. Since $X_i(\theta_i)$ is decreasing in θ_i and since $q(\theta, \varepsilon)$ is non-increasing in both ε (from Assumption 2) and θ, the second-order conditions about the two truthtelling constraints previously assumed to hold after equations (A.16) and (A.20) are satisfied. We are now ready to describe the first main result of our paper:

Theorem 1. *Together with the optimal decision rule $\boldsymbol{x}(\boldsymbol{\theta})$ from equations (2.17) and (2.18), as well as the transfers from equation (2.20), the quality policies $q(\theta, \varepsilon)$ that form part of the buyer optimal mechanism are given by the $q(\theta, \varepsilon)$ which solve:*

$$V'(q(\theta, \varepsilon)) = c_1(q(\theta, \varepsilon), \theta, \varepsilon) + \frac{F(\theta)}{f(\theta)}\left(c_{21}(q(\theta, \overline{\varepsilon}), \theta, \overline{\varepsilon}) - \frac{G(\varepsilon)}{g(\varepsilon)}c_{321}(q(\theta, \varepsilon), \theta, \varepsilon)\right) \tag{2.21}$$

$(\forall \theta, \varepsilon)$[5].

To make sense of this condition, note that the first-best socially optimal quality policies that maximise overall welfare and which the principal would prefer to implement without the presence of asymmetric information are instead given by the $q(\theta, \varepsilon)$ that solve:

$$V'(q(\theta, \varepsilon)) = c_1(q(\theta, \varepsilon), \theta, \varepsilon) \tag{2.23}$$

$(\forall \theta, \varepsilon)$. In order to limit the informational rent of the agents, the principal distorts the quality $q(\theta, \varepsilon)$ to be implemented by: $\frac{F(\theta)}{f(\theta)}\left(c_{21}(q(\theta, \overline{\varepsilon}), \theta, \overline{\varepsilon}) - \frac{G(\varepsilon)}{g(\varepsilon)}c_{321}(q(\theta, \varepsilon), \theta, \varepsilon)\right)$. Equation (2.21) can also be compared to the equation determining the buyer-optimal qualities used in Che

---

[5]We can equivalently express the first-order condition from equation (2.21) as:

$$V'(q(\theta, \varepsilon)) = c_1(q(\theta, \varepsilon), \theta, \varepsilon) + \frac{F(\theta)}{f(\theta)}\left(c_{21}(q(\theta, \underline{\varepsilon}), \theta, \underline{\varepsilon}) + \frac{1 - G(\varepsilon)}{g(\varepsilon)}c_{321}(q(\theta, \varepsilon), \theta, \varepsilon)\right) \tag{2.22}$$

$(\forall \theta, \varepsilon)$. To see this express the integral from equation (A.17) in terms of $v_i(\widetilde{\theta}_i, \theta_i, \boldsymbol{\theta}_{-i}, \underline{\varepsilon})$.

(1993) as discovered by Laffont and Tirole (1987), McAfee and J. McMillan (1987), and Riordan and Sappington (1987a):

$$V'(q(\theta)) = c_1(q(\theta), \theta) + \frac{F(\theta)}{f(\theta)} c_{21}(q(\theta), \theta) \tag{2.24}$$

($\forall \theta$). Here what determines the informational rent of the sellers is solely their type θ. In our adapted model with cost shocks, it also matters how the informational rent of an agent of type θ changes with the cost-shock ε. However, the effect of ε on the informational rent on the sellers is solely in how it changes how the informational rent changes with θ. The hazard rate of the θ distribution (through the ratio $F(\theta)/f(\theta)$) determines how the optimal qualities are adjusted downwards through the change in marginal cost for any ε. The hazard rate (through the ratio $G(\varepsilon)/g(\varepsilon)$) of the ε distribution determines how this adjustment changes with the cross derivative $c_{321}(q(\theta, \varepsilon), \theta, \varepsilon)$. The derivative $c_{31}(q(\theta, \varepsilon), \theta, \varepsilon)$ that determines how the informational rent changes with changes in the shock ε is not relevant to the optimal qualities which reflects that the shock is only relevant to the 'squeeze' in how it changes the squeeze as θ changes.

Crucial to how the squeeze changes with ε is the cross derivative c_{321}. When $c_{321} > 0$, the optimal mechanism squeezes lower realisations of the shock relatively more. On the contrary, when $c_{321} < 0$ we find that the squeeze for ε is decreasing in ε. As we would expect, if the distribution of the cost shock is degenerate, the squeeze of equation (2.21) becomes equivalent to the squeeze of Che from equation (2.24). Perhaps more interestingly, this is also the case when c_{321} is zero[6], since here $c_{21}(q, \theta, \underline{\varepsilon}) = c_{21}(q, \theta, \overline{\varepsilon})$. Unlike in Eső and Szentes (2007), when neither of these conditions is satisfied, i.e. when $\overline{\varepsilon} > \underline{\varepsilon}$ and $c_{21}(q, \theta, \underline{\varepsilon}) \neq c_{21}(q, \theta, \overline{\varepsilon})$, the agents' second stage private information affects their informational rents. In the optimal mechanism, the principal realises that the distribution of an agent's cost function is determined by their θ parameter (or their initial productivity) as well as the ε shock. Depending on the interactive term $c_{321}(q, \theta, \varepsilon)$, it is profitable for the principal to

---

[6]For example when $c(q, \theta, \varepsilon) = (\theta + \varepsilon) \cdot q$.

distinguish the make-up of the agents' costs in terms of θ and ε to optimally restrict the agents' informational rents.

2.4 Scoring Auctions with Renegotiation

Procurement contracts in practice are commonly awarded by some form of sealed bid first-price scoring auction. Here the bidders submit sealed bids with both price and scores for several quality attributes of a product or basket of products. The winner is selected as the bidder with the highest score and is asked to implement their bid of quality and price. Another in practice less commonly used scoring auction is the open bid ascending score auction. See Che (1993) for a discussion on different scoring auctions. In this section, we will introduce a scoring auction implementation game with built-in renegotiation, a model of scoring auctions with renegotiation (SAWR).

2.4.1 The SAWR Implementation Game

In this section a particular implementation game of **s**coring **a**uctions **w**ith **r**enegotiation for the model environment described in section 2.2 is introduced: the SAWR implementation game. This game is a stylised and streamlined timeline for procurement processes with cost shocks and renegotiation. Though stylised, it is designed to account for some of the critical characteristics of the process observed in collaboration with NHS Procurement Scotland. In summary, the SAWR implementation has the following timeline of events:

1. Agents realise their type θ.

2. The principal announces an auction with scoring rule $S(p, q)$ and whether the auction is first- or second-score.

3. The agents simultaneously bid a score S consisting of a price p and a quality q.

4. The agent with the highest score wins the auction at S_1 with the highest losing score being S_2. Without renegotiation, the agent would be asked to implement their bid p and q in a first-score auction. In a second-score auction the agent will be asked to implement their bid quality q at price p_2 such that $S(p_2, q) = S_2$.

5. Before any implementation, the winning agent realises ε.

6. The winning agent may choose a new p' and q' such that $S(p', q') \geq S_1$ at the original scoring rule in a first-score auction or $S(p', q') \geq S_2$ in a second-score auction.

7. The new p' and q' are now implemented.

Here 1 through to 4 mirror exactly the first-score and second-score auction as outlined by Che (1993). Point 5 crucially occurs after the winner is decided by the restriction imposed in the model setup of section 2.2. Point 6 details what will be referred to as the 'same-score restriction' for renegotiation. This restriction says that the seller can unilaterally decide a new price and quality after the realisation of the shock so long as the new price and quality match the original score S_1 they submitted to win the auction in point 3 and 4 (or S_2 in a second-score auction). This restriction serves two ends. First, by being unilateral, it is a very convenient and easy to model form of renegotiation. Second, it proxies for a 'fairness' restriction that describes the pertinent feature of many public procurement awards that the losing agents cannot be unfairly discriminated against by giving the winning agent post-allocation advantages (such as renegotiating to a lower score would be). In practice this is one of the most important considerations to public procurement institutions (such as NHS Procurement) when considering renegotiating the terms of a contract as the possibility of being sued by the losing agents is real and costly. Point 7 states that any implementation of price and quality only occurs after the renegotiation of price and quality. This is a stylisation that simplifies the analysis but does not qualitatively change the main results.

2.4.2 Equilibrium of SAWR

We restrict attention to quasilinear scoring rules that are linear in price in the same way as Che (1993):

Assumption 4. $S(p,q) = s(q) - p$, *where* $s(q) - c(q,\theta,\varepsilon)$ *has a unique interior maximum in* q *($\forall \theta, \varepsilon$) and* $s(\cdot)$ *is increasing at least for* $q \leq \mathrm{argmax}\, s(q) - c(q,\underline{\theta},\underline{\varepsilon})$.

These type of scoring rules are widespread in a procurement environment in practice. They can also be justified by the quasilinear preferences of our model. In equilibrium of the second stage of SAWR, an agent of type θ and shock ε thus chooses $q(\theta,\varepsilon)$ such that:

$$q(\theta,\varepsilon) = \underset{q}{\mathrm{argmax}}\, s(q) - c(q,\theta,\varepsilon). \tag{2.25}$$

In equilibrium of both first- and second score auctions the agent will choose the price $p(\theta,\varepsilon)$ in the second stage as:

$$p(\theta,\varepsilon) = s(q(\theta,\varepsilon)) - \overline{S}, \tag{2.26}$$

where $\overline{S}$ is the first stage score of the winning agent bid in the first-score auction and the score of the highest losing agent in the second-score auction. The quality choice in the second stage may thus be written in a value function:

$$S_o(\theta) = E_\varepsilon \left[\max_q \left\{ s(q) - c(q,\theta,\varepsilon) \right\} \right], \tag{2.27}$$

where $S_o(\theta)$ is the expected value of participation in the second stage of SAWR for an agent of type θ. The choice of the price p in the second stage is fully described by the choice of a score to bid in the first stage and the optimal qualities by equation (2.26). Now, the solution methodology of Che (1993) can be applied in expectation over the shock ε to reduce the first stage bidding of SAWR into a one-dimensional problem where agents bid according to their

expected value of the second stage. To this end, let as in Che (1993):

$$v \equiv S_o(\theta), \tag{2.28}$$

$$\widetilde{F}(v) \equiv 1 - F(S_o^{-1}(v)), \tag{2.29}$$

$$\text{and } b \equiv E_\varepsilon\Big[S(q(\theta,\varepsilon),p)\Big]. \tag{2.30}$$

$S_o(\theta)$ is strictly decreasing and therefore, its inverse exists. The bidding in the first stage of SAWR can then be described as agents choosing a bid b based on their productive potential v to maximise:

$$E_\varepsilon[\pi] = E_\varepsilon\Big[p - c(q(\theta,\varepsilon),\theta,\varepsilon)\Big] \cdot \text{Prob}\Big(\text{win}\Big| S(q(\theta,\varepsilon,p))\Big) \tag{2.31}$$

$$= [v - b] \cdot \widetilde{F}(b^{-1}(b))^{N-1}. \tag{2.32}$$

Leading to the following proposition:[7]

Proposition 2.4.1. *(i) A unique symmetric equilibrium of a quasilinear first-score SAWR mechanism is one in which each winning firm chooses:*

$$q(\theta,\varepsilon) = \underset{q}{\operatorname{argmax}}\; s(q) - c(q,\theta,\varepsilon) \tag{2.33}$$

and:

$$p(\theta,\varepsilon) = s(q(\theta,\varepsilon)) - \overline{S}^1(\theta), \tag{2.34}$$

where $\overline{S}^1(\theta)$ is the first stage score bid by the winning agent. In equilibrium of the first stage auction each agent bids:

$$\overline{S}(\theta) = E_\varepsilon\Big[s(q(\theta,\varepsilon))\Big] - E_\varepsilon\Big[c(q(\theta,\varepsilon))\Big] - \int_\theta^{\overline{\theta}} E_\varepsilon\Big[c_\theta(q(t,\varepsilon),t,\varepsilon)\Big]\Big(\frac{1 - F(t)}{1 - F(\theta)}\Big)^{N-1} dt. \tag{2.35}$$

(ii) A unique symmetric equilibrium of a second-score SAWR mechanism is one in which each winning firm chooses:

$$q(\theta,\varepsilon) = \underset{q}{\operatorname{argmax}}\; s(q) - c(q,\theta,\varepsilon) \tag{2.36}$$

---

[7]Analogous to proposition 2 in Che (1993).

and:

$$p(\theta, \varepsilon) = s(q(\theta, \varepsilon)) - \overline{S}^2(\theta), \tag{2.37}$$

where $\overline{S}^2(\theta)$ is the highest first stage score bid by a losing agent. In this equilibrium the agent has a dominant strategy in the first stage auction to bid:

$$\overline{S}(\theta) = E_\varepsilon\Big[s(q(\theta, \varepsilon)) - c(q(\theta, \varepsilon), \theta, \varepsilon)\Big]. \tag{2.38}$$

Proof. This is a straightforward extension of the equilibrium in proposition 2 of Che (1993). (i) and (ii) share the same second stage equilibrium. The first stage equilibrium of (i) follows from the usual equilibrium in first-price auctions as in Riley and Samuelson (1981). The first stage equilibrium of (ii) follows from Vickrey (1961). $\qquad\square$

The equilibrium expected profits of an agent of type θ in a first-score SAWR mechanism $E[\pi^{FS}]$ are thus:

$$E[\pi^{FS}] = E_\varepsilon[p - c(q(\theta, \varepsilon), \theta, \varepsilon)] \cdot \mathrm{P}(\text{win}|\theta) \tag{2.39}$$

$$= E_\varepsilon\Bigg[\int_\theta^{\overline{\theta}} c_\theta(q(t, \varepsilon), t, \varepsilon)\left(\frac{1 - F(t)}{1 - F(\theta)}\right)^{N-1} dt \cdot (1 - F(\theta))^{N-1}\Bigg] \tag{2.40}$$

$$= E_\varepsilon\Bigg[\int_\theta^{\overline{\theta}} c_\theta(q(t, \varepsilon), t, \varepsilon)(1 - F(t))^{N-1} dt\Bigg] = E[\pi^{SS}]. \tag{2.41}$$

Where $E[\pi^{SS}]$ are the expected equilibrium payoffs of a second score SAWR auction. Thus first and second score auctions are revenue equivalent for any general quasilinear scoring rule.

2.4.3 Does SAWR Implement the Optimal Mechanism?

Again we restrict attention to quasilinear scoring rules as proposed by Che (1993) and widely applied in practice.

Theorem 2. *When $c_{321} \neq 0$, no quasi-concave scoring rule $S(p, q) = s(q) - p$ exists with which the scoring auction with renegotiation (SAWR) mechanism described above can implement the optimal mechanism.*

Proof. This is proved by contradiction: assume a scoring rule $S(p,q)$ exists that implements the optimal mechanism. From the quasi-concavity we know that the upper contour set of $S(p,q)$ is convex. For any given score $\overline{S}$, the set of p and q such that $S(p,q) = \overline{S}$, can thus be mapped as a convex and non-decreasing function in the p,q space: an isoscore curve.

The indifference curves of a seller with preferences $\pi = p - c(q,\theta,\varepsilon)$ in the p,q space are strictly concave and strictly increasing given the assumptions on $c(q,\theta,\varepsilon)$. Thus, since $\pi_q < 0$ and $\pi_p > 0$, the bundle p^*, q^* that are the optimal way for the seller to achieve score $\overline{S}$ are unique and found where the seller's indifference curve is tangential to the isoscore curve of $\overline{S}$. See figure 2.1 for an illustration.
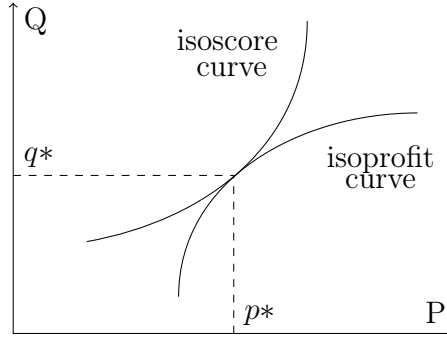
Figure 2.1 Isoscore and Isoprofit Curves

For two sellers $1, 2$ with $(\theta_1, \varepsilon_1)$ and $(\theta_2, \varepsilon_2)$ it's possible to get:

$$c(q, \theta_1, \varepsilon_1) = c(q, \theta_2, \varepsilon_2) \tag{2.42}$$

with $\theta_1 \neq \theta_2$, $\varepsilon_1 \neq \varepsilon_2$, and:

$$\frac{F(\theta_1)}{f(\theta_1)} \left(c_{21}(q(\theta_1, \overline{\varepsilon}), \theta_1, \overline{\varepsilon}) - \frac{G(\varepsilon_1)}{g(\varepsilon_1)} c_{321}(q(\theta_1, \varepsilon_1), \theta_1, \varepsilon_1) \right)$$
$$\neq \frac{F(\theta_2)}{f(\theta_2)} \left(c_{21}(q(\theta_2, \overline{\varepsilon}), \theta_2, \overline{\varepsilon}) - \frac{G(\varepsilon_2)}{g(\varepsilon_2)} c_{321}(q(\theta_2, \varepsilon_2), \theta_2, \varepsilon_2) \right) \tag{2.43}$$

(since $c_{321} \neq 0$). Given equation (2.42), at any score $\overline{S}$, the optimal p,q for both sellers are the same. However, from equation (2.21), equation (2.43) implies that

$$q(\theta_1, \varepsilon_1) \neq q(\theta_2, \varepsilon_2). \tag{2.44}$$

47

Thus a contradiction is reached. □

Scoring rules that are non-linear in price would need to reward a change in price differently at different levels of quality. We can convince ourselves that these would not help with this implementation problem by revisiting equation (2.43). The only way for sellers who have different θ and ε (and the same overall cost) to choose to procure the good at different qualities is for them to be indifferent between all these qualities. It is hard to conceive even a nonlinear scoring rule with which it is possible to ensure this condition while simultaneously ensuring that no sellers would want to procure the good at a quality meant for a seller with a different cost. However, a formal proof of this is beyond the scope of this analysis.

2.4.4 Naïve Scoring Auctions

In Che (1993), when $s(q) = V(q)$ the scoring auction implements the socially optimal 'first-best' level of qualities. This result extends to the SAWR mechanism:

Proposition 2.4.2. *If $s(q) = V(q)$ the SAWR implementation game implements the socially optimal quality policies q^* that solve:*

$$q^* = \underset{q}{\operatorname{argmax}} \, V(q) - c(q, \theta, \varepsilon) \tag{2.45}$$

($\forall \theta, \varepsilon$).

Proof. This is immediate from section 2.4.2 since $S(p, q) = V(q) - p$ is a quasilinear scoring rule. □

In the model without renegotiation, Che argues these qualities to be too high from the buyer's perspective as the naïve scoring rule 'fails to internalise the informational cost associated with increasing quality' (Che, 1993). However, an argument may be made that in government procurement, the government should at least to some extent care about the social welfare of its procurement mechanisms rather than exclusively expected buyer utility. Our

findings on the (sub)optimality of scoring auctions given renegotiation can be interpreted as furthering the case for 'naïve' auction and simplicity in procurement. First, note that a naïve scoring rule does not just implement the socially optimal qualities both in Che's analysis as well as in SAWR, but said scoring rule is exactly identical across the two models. The naïve scoring rule is thus robust to renegotiation and able to implement a socially optimal allocation even with unexpected renegotiation when the same score constraint is satisfied. Furthermore, in the limit as N grows large, the naïve SAWR approximates the optimal mechanism since the sellers' individual rent is decreasing in N. If in fact the government does allow for renegotiation of this form but does not foresee any shock, it may not take it into account in their scoring rule and decide to misrepresent their preferences in exactly the same way as Che suggests. In light of the following result, this can be a very costly mistake.

Proposition 2.4.3. *No quasilinear scoring rule in which the buyer's preferences are misrepresented can maintain the property of approximating the optimal mechanism as N grows large.*

Instead of a formal proof, we only provide the following heuristic to see this result. First, it is immediate from theorem 2 that no quality policies implementable by a quasilinear SAWR can squeeze θ and ε differently. This includes squeezing only θ or only ε. Second, from equation (2.21), note that in the optimal mechanism as N grows large the expected distortion of the quality policies of the eventual winner approaches zero in the limit since the expected first-order statistic of θ approaches $\underline{\theta}$. However, as N grows large, if the quality is distorted for both θ and ε in the same way, the limit of the distortion for the winner is never zero. Even though the first-order statistic for θ of the winning agent approaches $\underline{\theta}$, as there is only one remaining agent in the second stage the expectation of the shock ε does not change.

2.5 Conclusion

We extend a standard model of procurement to allow for the possibility of uncertainty about the true cost of production. Our main result is the description of the optimal mechanism in our multidimensional dynamic environment. This description allows us to make inferences about the performance of the popular scoring auction implementation game when we take into account the possibility of renegotiating in response to shocks. We show that the result of Che (1993) breaks down and scoring auctions are no longer optimal in this environment. Our analysis further suggests a policy recommendation for the design of government procurement auctions.

A crucial question in the design of public procurement processes is whether the government or public agency intending to procure a production process should only consider its own budget and preferences or whether a consideration to the sellers should also be made. Many argue the government's primary objective ought to be maximising social welfare. Laffont and Tirole (1987) view the objective function of the government as maximising the sum of its own utility and the bidders' utility but allow for some transaction costs to give some spiel to the government 'preferring its own preferences'. In both Che (1993) as well as our paper this is abstracted away, and the government is only concerned about maximising its own preferences. This is a key aspect of the design of scoring auctions. If the government does not care about the sellers' welfare, it should, as Che suggests, distort its preferences over non-price aspects of the procurement contract to implement as close to the optimal mechanism as possible. Here it accepts a socially worse outcome to reduce the informational rent of the sellers and increase its relative share of the overall gains from trade. On the other hand, if the government is equally interested in the sellers' welfare (and there are no transaction costs), the government is best off implementing a 'naïve' procurement auction in which it does not distort its own preferences and thus implements a socially optimal 'first-best' mechanism.

Our analysis can be used to further the case for transparency and simplicity in the design of procurement auctions by giving support to the argument for truthtelling and 'naïve' preference revelation in government procurement. We show that there are potential Pareto gains from renegotiation in multidimensional procurement environments and that revealing their true preferences over the non-price aspects of the good allows the procurer to realise these gains without requiring any distributional knowledge of the shock distribution. The naïve scoring auction here has the desirable property of being able to make use of the same scoring rule to allocate the contract as well as to use in a renegotiation phase, should the need for one arise. In this sense, it is robust to renegotiation and can ensure that the scoring auction implements a dynamic contract that is robust to the potential of double deviations and satisfies a government's commitment to fairness over favouritism in the procurement process. In contrast, scoring auctions with preference distortions aimed at reducing the sellers' rents, on top of being computationally demanding and requiring strong informational assumptions to be implemented correctly, do not satisfy this property of being robust to renegotiation. They implement an allocation that is undesirable for both a government concerned about its own utility and a government maximising social welfare.

Appendix

A.1 Proof of Lemma 2.3.1

Two difficulties arise when showing the revelation principle holds in this environment. First, the mechanism design problem is dynamic, with two stages and sequential generation of information. Second, the mechanism design problem is restricted by enforcing a winner of the contract before the realisation of the shock.

A general mechanism Γ in our model setting defines: **(i)** the strategies available to each agent $\Phi_1, \cdots, \Phi_N$ and **(ii)** an outcome function g for each realisation of strategies by the agents. A realisation $\phi_i(\theta_i) \in \Phi_i$ of a strategy for an agent $i \in \{1, \cdots, N\}$ dictates an action for the agent at every possible node of the game tree. The strategy thus determines the agent's actions before the allocation decision and after when the shock ε has been realised. The outcome function g defines **(a)** an allocation decision x_i ($\forall i$), **(b)** transfers t_i to the agent, and finally **(c)** a quality q_i to be implemented ($\forall i$) if the agent is chosen by the allocation rule x_i. Here g is a function of the agents' realisations of strategies. Thus g defines $\mathbf{x} \equiv (x_1, \cdots, x_N)$, $\mathbf{t} \equiv (t_1, \cdots, t_N)$, and $\mathbf{q} \equiv (q_1, ..., q_N)$ for all $\boldsymbol{\phi}(\boldsymbol{\theta}) \equiv (\phi_1(\theta_1), \cdots, \phi_N(\theta_N)) \in (\Phi_1, \cdots, \Phi_N) \equiv \boldsymbol{\Phi}$.

The principal can design both the strategy space $\boldsymbol{\Phi}$ as well as the function g. In an equilibrium of Γ, agents use equilibrium strategies $\phi_1^*(\theta_1), \cdots, \phi_N^*(\theta_N) \equiv \boldsymbol{\phi}^*(\boldsymbol{\theta})$ resulting in equilibrium outcome $g(\boldsymbol{\phi}^*(\boldsymbol{\theta}))$. To characterise the equilibrium strategies with the perfect Bayesian equilibrium concept, a division is helpful. Let $\phi_i^1(\theta_i)$ denote the part of the strategy

$\phi_i(\theta_i)$ that comprises of all decisions in the game tree taken *before* the realisation of ε ($\forall i$). Similarly, let $\phi_i^2(\theta_i, \varepsilon | \phi_i^1(\theta_i))$ denote the part of the strategy $\phi_i(\theta_i)$ that comprises of all decisions in the game tree taken *after* the realisation of ε ($\forall i$) conditional on the agent's type, the realised shock, and the decisions taken before the shock. $\phi_i(\theta_i)$ is thus described by $\phi_i^1(\theta_i)$ as well as $\phi_i^2(\theta_i, \varepsilon | \phi_i^1(\theta_i))$.

The aim is to show that any perfect Bayesian equilibrium outcome of any general mechanism Γ can be implemented in perfect Bayesian equilibrium by a direct mechanism Γ^{DM} in which the agents simultaneously and truthfully reveal their type θ to the principal and the agent winning the allocation decision reveals ε directly and truthfully to the principal. The first thing to note here is that in the direct mechanism, the winning agent does not observe the other agents' type reports (they only observe that they win). We have not, however, ruled out general mechanisms in which the agent can observe part of or all of the other agents' actions and make inferences about their strategies and types. The expectation over $\boldsymbol{\theta}_{-i}$ in the second stage is implicitly conditional on whatever was observed in the first stage and may thus be different across the direct mechanism and some general mechanisms. To circumvent this informational disparity in the conditional expectations, we require the principal to be able to reveal to the winning agent information about the other agents' types before the beginning of the second stage (and to commit to doing so truthfully)[1].

A direct mechanism is thus a constraint on the general mechanism such that $\Phi_i^1 = \Theta$ ($\forall i$) and $\Phi_i^2 = \mathfrak{E}$ ($\forall i$). It specifies further a message $\widehat{I}$ about the other agents' types from the principal to the remaining agent in the second stage of the game tree. In order for Γ^{DM} to implement the same equilibrium outcome as Γ, we adjust the outcome function g^{DM} in

---

[1]Ultimately, given lemma 2.3.2 and corollary 2.3.2.1 of section 2.3, we find that under independent private valuations this revelation is not relevant for the optimal direct mechanism. However, to show that restricting attention to direct mechanisms does not restrict the principal's optimisation problem over all possible mechanisms, we first need to show that any general mechanism can be implemented with a direct one. Ruling out by assumption any mechanism in which the agents learn about each others' types is thus not obviously non-restrictive to the principal.

such a way that any vector of reports $\widehat{\boldsymbol{\theta}}$ jointly with any report $\widehat{\varepsilon}$ by the remaining agent j results in the same outcome g as the realisation $\boldsymbol{\phi}^*(\widehat{\boldsymbol{\theta}})$ would have in Γ.

Lemma A.1.1. The Revelation Principle *Suppose that there exists a mechanism* $\Gamma = (\Phi_1^1, \cdots, \Phi_N^1, \Phi_1^2, \cdots, \Phi_N^2, g(\cdot))$ *that implements[2] an outcome function f in perfect Bayesian Nash equilibrium. Then f is truthfully implementable in perfect Bayesian Nash equilibrium by a direct mechanism Γ^{DM} as defined above.*

Proof. This proof is adapted from proposition 23.D.1 in Mas-Colell, Whinston, Green, et al. (1995).

As will become clear below, we require the principal to commit to revealing to any winning agent i (with first stage report $\widehat{\theta}_i$) before stage 2 of the DM whatever information the agent would have observed in stage 1 of the general mechanism had he followed $\phi_i^{1*}(\widehat{\theta}_i)$. This information is available to the principal before the second stage as she observes the vector of type reports and it is otherwise not available to the winning agent in the direct mechanism as he does not observe it.

If Γ implements $f(\cdot)$ in perfect Bayesian equilibrium, then there exists a profile of equilibrium strategies $\boldsymbol{\phi}^{*1}(\boldsymbol{\theta}), \boldsymbol{\phi}^{*2}(\boldsymbol{\theta}, \varepsilon)$ such that $g(\boldsymbol{\phi}^{*1}(\boldsymbol{\theta}), \boldsymbol{\phi}^{*2}(\boldsymbol{\theta}, \varepsilon)) = f(\boldsymbol{\theta}, \varepsilon)$. Thus the equilibrium first stage strategies have to satisfy:

$$\phi_i^{1*}(\theta_i) \in \underset{\widetilde{\phi}_i^1(\theta_i)}{\operatorname{argmax}} E_{\varepsilon, \boldsymbol{\theta}_{-i}} \left[\pi_i \left(g\left(\widetilde{\phi}_i^1(\theta_i), \widehat{\phi}_i^2(\theta_i, \varepsilon | \widetilde{\phi}_i^1(\theta_i)), \boldsymbol{\phi}_{-i}^*(\boldsymbol{\theta}_{-i}) \right), \theta_i, \varepsilon \right) \Big| I_i^1 \left(\widetilde{\phi}_i^1(\theta_i) \right) \right] \quad \text{(A.1)}$$

$(\forall i, \theta_i)$, here $I_i(\phi)$ represents the information[3] the agent i learns about the other agents' types $\boldsymbol{\theta}_{-i}$ when following strategy ϕ. $\widehat{\phi}_i^2(\theta_i, \varepsilon | \widetilde{\phi}_i^1(\theta_i))$ is an optimal second stage response defined by:

$$\widehat{\phi}_i^2 \in \underset{\widetilde{\phi}_i^2(\theta_i, \varepsilon | \widetilde{\phi}_i^1(\theta_i))}{\operatorname{argmax}} E_{\boldsymbol{\theta}_{-i}} \left[\pi_i \left(g\left(\widetilde{\phi}_i^1(\theta_i), \widetilde{\phi}_i^2(\theta_i, \varepsilon | \widetilde{\phi}_i^1(\theta_i)), \boldsymbol{\phi}_{-i}^*(\boldsymbol{\theta}_{-i}) \right), \theta_i, \varepsilon \right) \Big| I_i \left(\widetilde{\phi}_i^1(\theta_i) \right) \right] \quad \text{(A.2)}$$

---

[2]We are not assuming uniqueness here.

[3]Given its role in the analysis, it is not necessary to describe the functional form of this message explicitly.

$(\forall i, \theta_i, \varepsilon, \widetilde{\phi}_i^1(\theta_i))$. In words, the equilibrium first stage strategies have to be optimal in expectation over ε given subgame perfection in the second stage. This implies that the equilibrium second stage strategies have to satisfy:

$$\phi_i^{2*} \in \operatorname*{argmax}_{\widetilde{\phi}_i^2(\theta_i, \varepsilon | \phi_i^{1*}(\theta_i))} E_{\boldsymbol{\theta}_{-i}} \left[\pi_i \left(g \left(\phi_i^{1*}(\theta_i), \widetilde{\phi}_i^2(\theta_i, \varepsilon | \phi_i^{1*}(\theta_i)), \boldsymbol{\phi}_{-i}^*(\boldsymbol{\theta}_{-i}) \right), \theta_i, \varepsilon \right) \Big| I_i \left(\phi_i^{1*}(\theta_i) \right) \right] \quad \text{(A.3)}$$

$(\forall i, \theta_i, \varepsilon, \phi_i^{1*}(\theta_i))$. In words, given equation (A.1) and (A.2) and using the perfect Bayesian equilibrium concept, the second stage equilibrium strategies only have to be optimal at the optimal first stage strategy ϕ_i^{1*} for the bundle ϕ_i^{1*}, ϕ_i^{2*} to be an equilibrium strategy. Now, for $f(\cdot)$ to be implementable in truthtelling equilibrium it has to be that:

$$\theta_i \in \operatorname*{argmax}_{\widetilde{\theta}_i} E_{\varepsilon, \boldsymbol{\theta}_{-i}} \left[\pi_i \left(f \left(\widetilde{\theta}_i, \boldsymbol{\theta}_{-i}, \widehat{\varepsilon} \right), \theta_i, \varepsilon \right) \Big| \widehat{I}_i \left(\widetilde{\theta}_i \right) \right] \quad \text{(A.4)}$$

$(\forall i, \theta_i)$, Here $\widehat{I}_i \left(\widetilde{\theta}_i \right)$ is the information revealed to agent i if they win the allocation of the second stage given type report $\widetilde{\theta}_i$ in the first stage of the DM. $\widehat{\varepsilon}$ is the best response defined by:

$$\widehat{\varepsilon} \in \operatorname*{argmax}_{\widetilde{\varepsilon}} E_{\boldsymbol{\theta}_{-i}} \left[\pi_i \left(f \left(\widetilde{\theta}_i, \boldsymbol{\theta}_{-i}, \widetilde{\varepsilon} \right), \theta_i, \varepsilon \right) \Big| \widehat{I}_i \left(\widetilde{\theta}_i \right) \right] \quad \text{(A.5)}$$

$(\forall i, \theta_i, \widetilde{\theta}_i, \varepsilon)$. Furthermore, it has to hold that:

$$\varepsilon \in \operatorname*{argmax}_{\widetilde{\varepsilon}} E_{\boldsymbol{\theta}_{-i}} \left[\pi_i \left(f \left(\theta_i, \boldsymbol{\theta}_{-i}, \widetilde{\varepsilon} \right), \theta_i, \varepsilon \right) \Big| \widehat{I}_i \left(\theta_i \right) \right] \quad \text{(A.6)}$$

$(\forall i, \theta_i, \varepsilon)$. Since $f(\boldsymbol{\theta}, \varepsilon) = g(\boldsymbol{\phi}^{*1}(\boldsymbol{\theta}), \boldsymbol{\phi}^{*2}(\boldsymbol{\theta}, \varepsilon))$, the first step is to confirm that:

$$\theta_i \in \operatorname*{argmax}_{\widetilde{\theta}_i} E_{\varepsilon, \boldsymbol{\theta}_{-i}} \left[\pi_i \left(g \left(\phi_i^{1*}(\widetilde{\theta}_i), \phi_i^{2*}(\widetilde{\theta}_i, \widehat{\varepsilon} | \phi_i^{1*}(\widetilde{\theta}_i)), \boldsymbol{\phi}_{-i}^*(\boldsymbol{\theta}_{-i}) \right), \theta_i, \varepsilon \right) \Big| \widehat{I}_i \left(\widetilde{\theta}_i \right) \right] \quad \text{(A.7)}$$

$(\forall i, \theta_i)$, with:

$$\widehat{\varepsilon} \in \operatorname*{argmax}_{\widetilde{\varepsilon}} E_{\boldsymbol{\theta}_{-i}} \left[\pi_i \left(g \left(\phi_i^{1*}(\widetilde{\theta}_i), \phi_i^{2*}(\widetilde{\theta}_i, \widetilde{\varepsilon} | \phi_i^{1*}(\widetilde{\theta}_i)), \boldsymbol{\phi}_{-i}^*(\boldsymbol{\theta}_{-i}) \right), \theta_i, \varepsilon \right) \Big| \widehat{I}_i \left(\widetilde{\theta}_i \right) \right] \quad \text{(A.8)}$$

$(\forall i, \theta_i, \widetilde{\theta}_i, \varepsilon)$.

If the principal reveals to agent i with report $\widetilde{\theta}_i$ the information they would acquire following the strategy $\phi_i^{1*}(\widetilde{\theta}_i)$, i.e. if $\widehat{I}_i \left(\widetilde{\theta}_i \right) = I_i \left(\phi_i^{1*}(\widetilde{\theta}_i) \right)$ the conditional expectations over

$\boldsymbol{\theta}_{-i}$ in equation (A.2) and (A.5) are identical. Thus, since $\phi_i^{2*}(\widetilde{\theta}_i, \widetilde{\varepsilon}|\phi_i^{1*}(\widetilde{\theta}_i))$ is available as a deviation $\widetilde{\phi}_i^2(\theta_i, \varepsilon|\widetilde{\phi}_i^1(\theta_i))$ in equation (A.2) and since $\phi_i^{1*}(\widetilde{\theta}_i)$ is available as a deviation $\widetilde{\phi}_i^1(\theta_i)$ in equation (A.1), $\phi_i^{1*}(\theta_i)$ being an equilibrium strategy guarantees that (A.7) with (A.8) and thus (A.4) with (A.5) are satisfied. To confirm that $f(\cdot)$ is implementable in truthtelling equilibrium in this case, it now only needs to be checked whether:

$$\varepsilon \in \underset{\widetilde{\varepsilon}}{\operatorname{argmax}} E_{\boldsymbol{\theta}_{-i}} \left[\pi_i \left(g \left(\phi_i^{1*}(\theta_i), \phi_i^{2*}(\theta_i, \widetilde{\varepsilon}|\phi_i^{1*}(\theta_i)), \boldsymbol{\phi}_{-i}^*(\boldsymbol{\theta}_{-i}) \right), \theta_i, \varepsilon \right) \Big| \widehat{I}_i(\theta_i) \right] \qquad \text{(A.9)}$$

$(\forall i, \theta_i, \varepsilon)$. This is relatively straightforward. First, note that the expectations over $\boldsymbol{\theta}_{-i}$ in equation (A.3) and (A.6) are identical when $\widehat{I}_i(\theta_i) = I_i(\phi_i^{1*}(\theta_i))$. Now, $\phi_i^{2*}(\theta_i, \widetilde{\varepsilon}|\phi_i^{1*}(\theta_i))$ is available as a deviation $\widetilde{\phi}_i^2(\theta_i, \varepsilon|\phi_i^{1*}(\theta_i))$ in equation (A.3). Thus $\phi_i^{2*}(\theta_i, \varepsilon|\phi_i^{1*}(\theta_i))$ being an equilibrium strategy guarantees that (A.9) and thus (A.6) are satisfied since the expectations over $\boldsymbol{\theta}_{-i}$ are the same in equations (A.3) and (A.9) given the discussion above.

$\square$

A.2 Proof of Lemma 2.3.2

This is proven by contradiction. Suppose there is some agent i for whom the optimal quality policy denoted q_i depends not only on their own cost type θ_i as well as the shock ε, but also on the vector $\boldsymbol{\theta}_{-i}$ of the $N-1$ other agents' type realisations. For this agent i, it has to be such that there are at least two different realisations of $\boldsymbol{\theta}_{-i}$ denoted $\boldsymbol{\theta}_{-i}^1$ and $\boldsymbol{\theta}_{-i}^2$ such that the quality to be implemented by agent i is different between $\boldsymbol{\theta}_{-i}^1$ and $\boldsymbol{\theta}_{-i}^2$, but the agent i wins the allocation at both $\boldsymbol{\theta}_{-i}^1$ and $\boldsymbol{\theta}_{-i}^2$. Denote the two different qualities q_i^1 and q_i^2, respectively.

If the principal were to offer a different mechanism to the N agents, in which at both $\boldsymbol{\theta}_{-i}^1$ and $\boldsymbol{\theta}_{-i}^2$ agent i is asked to implement $E[q]$ such that[4]:

$$q(\theta_i, \boldsymbol{\theta}_{-i}^1, \varepsilon) = q(\theta_i, \boldsymbol{\theta}_{-i}^2, \varepsilon) = \frac{Pr(\boldsymbol{\theta}_{-i}^1)}{Pr(\boldsymbol{\theta}_{-i}^1) + Pr(\boldsymbol{\theta}_{-i}^2)} \cdot q_i^1 + \frac{Pr(\boldsymbol{\theta}_{-i}^2)}{Pr(\boldsymbol{\theta}_{-i}^1) + Pr(\boldsymbol{\theta}_{-i}^2)} \cdot q_i^2 \equiv E[q_i] \quad \text{(A.10)}$$

---

[4]Here $Pr(\cdot)$ denotes the likelihood of a particular vector $\boldsymbol{\theta}_{-i}$.

the principal would make an expected profit. To see this, notice that if transfers and allocation decisions in the new mechanism remain unchanged, due to the concavity of the principal's preferences over quality, the principal increases their expected utility by changing the mechanism in this way. Mathematically:

$$V(E[q]) > \frac{Pr(\boldsymbol{\theta}_{-i}^1)}{Pr(\boldsymbol{\theta}_{-i}^1) + Pr(\boldsymbol{\theta}_{-i}^2)} \cdot V(q_i^1) + \frac{Pr(\boldsymbol{\theta}_{-i}^2)}{Pr(\boldsymbol{\theta}_{-i}^1) + Pr(\boldsymbol{\theta}_{-i}^2)} \cdot V(q_i^2) \tag{A.11}$$

by Jensen's inequality and given $V'(q) > 0$ and $V''(q) < 0$. Furthermore, from the convexity of the agents' cost:

$$c(E[q], \theta_i, \varepsilon) < \frac{Pr(\boldsymbol{\theta}_{-i}^1)}{Pr(\boldsymbol{\theta}_{-i}^1) + Pr(\boldsymbol{\theta}_{-i}^2)} \cdot c(q_i^1, \theta_i, \varepsilon) + \frac{Pr(\boldsymbol{\theta}_{-i}^2)}{Pr(\boldsymbol{\theta}_{-i}^1) + Pr(\boldsymbol{\theta}_{-i}^2)} \cdot c(q_i^2, \theta_i, \varepsilon) \tag{A.12}$$

by Jensen's inequality and given $c_1(q, \theta_i, \varepsilon) > 0$ and $c_{11}(q, \theta_i, \varepsilon) > 0$ using the independence of values assumed for the cost function. The expected cost of implementing $E[q_i]$ is thus lower than the expected cost of implementing q_i^1 and q_i^2. Thus replacing q_i^1 and q_i^2 by $E[q_i]$ improves the optimisation programme. This contradicts the earlier assertion that q_i^1 and q_i^2 are optimal.

A.3   Proof of Lemma 2.3.3

From equation (2.4) the expected value of winning the allocation decision for agent i of type θ_i when realising shock ε and reporting type $\widetilde{\theta}_i$ in the first stage is:

$$v_i(\widetilde{\theta}_i, \theta_i, \varepsilon) = \max_{\widetilde{\varepsilon}} E_{\boldsymbol{\theta}_{-i}} \left[t_i^2(\widetilde{\theta}_i, \widetilde{\varepsilon}) - c(q_i(\widetilde{\theta}_i, \widetilde{\varepsilon}), \theta_i, \varepsilon) \right] \tag{A.13}$$

$$= \max_{\widetilde{\varepsilon}} \left\{ t_i^2(\widetilde{\theta}_i, \widetilde{\varepsilon}) - c(q_i(\widetilde{\theta}_i, \widetilde{\varepsilon}), \theta_i, \varepsilon) \right\} \tag{A.14}$$

Thus letting:

$$\frac{\partial v_i(\widetilde{\theta}_i, \theta_i, \varepsilon)}{\partial \varepsilon} = -c_3(q_i(\widetilde{\theta}_i, \widetilde{\varepsilon}), \theta_i, \varepsilon) \Big|_{\widetilde{\varepsilon} = \varepsilon} \tag{A.15}$$

$$= -c_3(q_i(\widetilde{\theta}_i, \varepsilon), \theta_i, \varepsilon) \tag{A.16}$$

$(\forall i, \widetilde{\theta}_i, \theta_i, \varepsilon$ such that $\widetilde{\theta}_i = \theta_i)$ is a necessary condition for the truthtelling constraint to be satisfied (by the envelope theorem)[5]. This condition is sufficient if $q(\theta_i, \varepsilon)$ is non-increasing in ε. This will be assumed for the time being and can later be shown to hold under assumption 2. Since condition 2.5 only has to hold at $\widetilde{\theta}_i = \theta_i$, truthtelling about ε in the second stage is enforced when the derivative of v fulfils this condition at $\widetilde{\theta}_i = \theta_i$. Integrating both sides of (A.16):

$$v_i(\widetilde{\theta}_i, \theta_i, \varepsilon) = \int_{\varepsilon}^{\overline{\varepsilon}} c_3(q_i(\widetilde{\theta}_i, s), \theta_i, s)ds + v_i(\widetilde{\theta}_i, \theta_i, \overline{\varepsilon}) \tag{A.17}$$

gives the value of the second stage $(\forall i, \widetilde{\theta}_i, \theta_i, \varepsilon$ such that $\widetilde{\theta}_i = \theta_i)$ given constraint (2.5).

A.4   Proof of Lemma 2.3.4

First note that:

$$W_i(\theta_i) = \max_{\widetilde{\theta}_i} E_{\varepsilon, \boldsymbol{\theta}_{-i}} \left[v_i(\widetilde{\theta}_i, \theta_i, \varepsilon) x_i(\widetilde{\theta}_i, \boldsymbol{\theta}_{-i}) + t_i^1(\widetilde{\theta}_i, \boldsymbol{\theta}_{-i}) \right] \tag{A.18}$$

$(\forall i, \theta_i)$. From the truthtelling constraint in equation (2.3) to get:

$$W_i(\theta_i) = E_{\varepsilon, \boldsymbol{\theta}_{-i}} \left[v_i(\widetilde{\theta}_i, \theta_i, \varepsilon) x_i(\widetilde{\theta}_i, \boldsymbol{\theta}_{-i}) + t_i^1(\widetilde{\theta}_i, \boldsymbol{\theta}_{-i}) \right] \Big|_{\widetilde{\theta}_i = \theta_i} \tag{A.19}$$

$(\forall i, \theta_i)$. Thus letting:

$$\frac{dW_i(\theta_i)}{d\theta_i} = E_{\varepsilon, \boldsymbol{\theta}_{-i}} \left[\frac{\partial v_i(\widetilde{\theta}_i, \theta_i, \varepsilon)}{\partial \theta_i} x_i(\widetilde{\theta}_i, \boldsymbol{\theta}_{-i}) \right] \Big|_{\widetilde{\theta}_i = \theta_i} \tag{A.20}$$

$(\forall \theta_i)$ is a necessary condition for truthtelling about θ (by the envelope theorem)[6]. When $E_{\boldsymbol{\theta}_{-i}}[x_i(\boldsymbol{\theta})]$ is non-increasing in θ_i, and $q_i(\theta_i, \varepsilon)$ is non-increasing in θ_i it is a sufficient condition for truthtelling. We will later show these to be satisfied under assumption 2 and proceed as though they are for now. Notice that:

$$E_{\varepsilon} \left[v_i(\widetilde{\theta}_i, \theta_i, \varepsilon) \right] \Big|_{\widetilde{\theta}_i = \theta_i} = E_{\varepsilon} \left[\int_{\varepsilon}^{\overline{\varepsilon}} c_3(q_i(\widetilde{\theta}_i, s), \theta_i, s)ds + v_i(\widetilde{\theta}_i, \theta_i, \overline{\varepsilon}) \right] \Big|_{\widetilde{\theta}_i = \theta_i} \tag{A.21}$$

---

[5]The differentiability of v stems from the differentiability of c.

[6]W is differentiable because v is differentiable; v because c is.

$(\forall i, \theta_i)$. Here we could substitute for equation (2.6) in the last line since we are evaluating at $\widetilde{\theta}_i = \theta_i$. Taking the derivative w.r.t. θ_i on both sides we get:

$$\frac{\partial E_\varepsilon \left[v_i(\widetilde{\theta}_i, \theta_i, \varepsilon) \right]}{\partial \theta_i} \Bigg|_{\widetilde{\theta}_i=\theta_i} = \frac{\partial E_\varepsilon \left[\int_\varepsilon^{\overline{\varepsilon}} c_3(q_i(\widetilde{\theta}_i, s), \theta_i, s) ds + v_i(\widetilde{\theta}_i, \theta_i, \overline{\varepsilon}) \right]}{\partial \theta_i} \Bigg|_{\widetilde{\theta}_i=\theta_i} \tag{A.22}$$

$$= E_\varepsilon \left[\int_\varepsilon^{\overline{\varepsilon}} c_{32}(q_i(\widetilde{\theta}_i, s), \theta_i, s) ds + \frac{\partial v_i(\widetilde{\theta}_i, \theta_i, \overline{\varepsilon})}{\partial \theta_i} \right] \Bigg|_{\widetilde{\theta}_i=\theta_i} \tag{A.23}$$

$(\forall i, \theta_i)$. From equation (A.13) we can get:

$$\frac{\partial v_i(\widetilde{\theta}_i, \theta_i, \overline{\varepsilon})}{\partial \theta_i} = - -c_2(q_i(\widetilde{\theta}_i, \widetilde{\varepsilon}), \theta_i, \overline{\varepsilon}) \Big|_{\widetilde{\varepsilon}=\varepsilon^*} \tag{A.24}$$

$(\forall i, \theta_i)$ using the envelope theorem. Here ε^* represents the optimal choice for ε given the state variables. Therefore:

$$\frac{\partial \left[v_i(\widetilde{\theta}_i, \theta_i, \overline{\varepsilon}) \right]}{\partial \theta_i} \Bigg|_{\widetilde{\theta}_i=\theta_i} = -c_2(q_i(\widetilde{\theta}_i, \widetilde{\varepsilon}), \theta_i, \overline{\varepsilon}) \Big|_{\substack{\widetilde{\varepsilon}=\varepsilon^* \\ \theta_i=\theta_i}} \tag{A.25}$$

$(\forall i, \theta_i)$. Given second stage truthtelling, $\varepsilon^* = \varepsilon$ when $\widetilde{\theta}_i = \theta_i$, thus:

$$\frac{\partial v_i(\widetilde{\theta}_i, \theta_i, \overline{\varepsilon})}{\partial \theta_i} \Bigg|_{\widetilde{\theta}_i=\theta_i} = -c_2(q_i(\widetilde{\theta}_i, \overline{\varepsilon}), \theta_i, \overline{\varepsilon}) \Big|_{\widetilde{\theta}_i=\theta_i} \tag{A.26}$$

$(\forall i, \theta_i)$. Substituting into equation (A.20) we get:

$$\frac{dW_i(\theta_i)}{d\theta_i} = E_{\varepsilon,\boldsymbol{\theta}_{-i}} \left[\left(\int_\varepsilon^{\overline{\varepsilon}} c_{32}(q_i(\widetilde{\theta}_i, s), \theta_i, s) ds - c_2(q_i(\widetilde{\theta}_i, \overline{\varepsilon}), \theta_i, \overline{\varepsilon}) \right) x_i(\widetilde{\theta}_i, \boldsymbol{\theta}_{-i}) \right] \Bigg|_{\widetilde{\theta}_i=\theta_i}$$

$$= E_{\varepsilon,\boldsymbol{\theta}_{-i}} \left[\left(\int_\varepsilon^{\overline{\varepsilon}} c_{32}(q_i(\theta_i, s), \theta_i, s) ds - c_2(q_i(\theta_i, \overline{\varepsilon}), \theta_i, \overline{\varepsilon}) \right) x_i(\theta_i, \boldsymbol{\theta}_{-i}) \right] \tag{A.27}$$

$(\forall i, \theta_i)$. From (A.27), $W_i(\theta_i)$ is non-increasing in θ_i since $c_{32} \geq 0$ and $c_2 > 0$. Thus for the participation constraint (2.2) to be satisfied it is enough for it to be satisfied for a agent of type $\overline{\theta}$, i.e. for:

$$W_i(\overline{\theta}) = 0 \tag{A.28}$$

to hold $(\forall i)$. Thus:

$$W_i(\theta_i) = E_{\varepsilon,\boldsymbol{\theta}_{-i}} \left[\int_{\theta_i}^{\overline{\theta}} \left(c_2(q_i(t, \overline{\varepsilon}), t, \overline{\varepsilon}) - \int_\varepsilon^{\overline{\varepsilon}} c_{32}(q_i(t, s), t, s) ds \right) x_i(t, \boldsymbol{\theta}_{-i}) dt \right] \tag{A.29}$$

$(\forall i, \theta_i)$, which concludes the proof.

60

Chapter Three

OPTIMAL SCORING AUCTIONS WITH RENEGOTIATION

3.1 Introduction

Mechanism design adds an element of reverse engineering to game-theoretic applications. In strategic settings, when agents hold private information, we can use game theory to describe the equilibrium outcomes of the agents' interactions. For example, in the context of first-price auctions, we can describe agents' bidding strategies as best response functions to the other agents' strategies. Equilibrium is reached where all agents are best responding to one another's best response. This analysis allows us to predict the expected revenue for the auctioneer, or an optimal strategy for a bidder in the auction. However, often the auctioneer is not constrained to selling the good through first-price auctions. There might be a large number of possible implementation games she can choose from when selling a good. Asides from a first-price auction, there are second-price auctions, double auctions, all-pay auctions, and many others. On top of that, she could negotiate with one of the bidders, hold a contest, or even just set a price and sell the good to a randomly selected bidder. The possibilities are endless. Mechanism design allows us to abstract away from the specifics of the implementation game and describe the properties of the best possible allocation that

can be implemented in equilibrium by a general mechanism in a given model environment. As such, it allows us to evaluate our chosen implementation games by comparing their equilibrium allocations with that of the best possible allocation.

In environments with complex type-spaces and dynamics, the description of the optimal mechanism becomes more and more complex. With this added complexity, identifying a realistic implementation game of the optimal mechanism becomes increasingly difficult. The informational assumptions may be too strong, or the computational demands too high to expect the optimal mechanism to find real-world application. Moreover, there may be implementation games that are appealing even if they are unable to implement the allocation of the theoretically optimal mechanism. For example, in the context of government procurement, goods are commonly contracted over various non-price elements and long time frames in which unforeseen dynamics can quickly arise. In chapter two of this thesis, we show that the prevalent procurement mechanism, scoring auctions with renegotiation in response to unforeseen changes, is not able to implement the optimal mechanism if shocks to production cost can arise after the initial auction. Despite this, scoring auctions remain of interest, primarily due to their common use in practice. However, we are only able to give a negative result about scoring auctions with renegotiation. We do not learn what the best scoring rule for an auction with renegotiation looks like and how much is lost compared to the optimal mechanism.

In this chapter, we will introduce a simplified environment compared to chapter 2. We will restrict attention to a discrete type-space and cost-shock. This simplification has several benefits. Firstly, reproducing the key result of chapter 2, that the scoring auction with renegotiation is not optimal, in a discrete setting provides some helpful intuition to guide our understanding of this finding. Secondly, the discrete environment allows us to conceptualise and analytically describe the optimal scoring auction. That is, given that scoring auctions are sub-optimal, we discover the optimal mechanism among the sub-class of mechanisms that are implementable through the scoring auction with renegotiation implementation game. This

analysis is, to the best of our understanding, a new and unique contribution to the theory of mechanism design. Finally, having described the optimal scoring auctions, we can, using a simple numerical example, develop an understanding of how much, relative to the optimal mechanism, is lost by a government's commitment to the scoring auction.

The outline of this paper is as follows. Section 3.2 describes the general model setup. Section 3.3 outlines a motivating numerical example which we use to guide the analysis throughout the paper intuitively. In section 3.4, we derive the solution to the optimal mechanism problem and compare it to previous literature. Section 3.5 introduces a scoring auction with renegotiation mechanism formally. In section 3.6 we describe the properties of the optimal scoring auction with renegotiation using a restricted mechanism design approach. Section 3.7 offers some concluding remarks.

3.2 The General Nx2 Model

This model setup takes the model from chapter 2 to a discrete type-space. There are $N \in \mathbb{R}^+$ risk-neutral sellers (agents) indexed i trying to sell to a single risk-neutral buyer (principal) a single indivisible good. The agents and the principal have preferences over two-dimensions of the good to be contracted: its price p and its quality q. The sellers differ in their cost of producing quality $c(q, \theta_i, \varepsilon)$ by a cost parameter $\theta_i = \{\bar{\theta}, \underline{\theta}\}$ ($\forall i$) and a common shock parameter $\varepsilon = \{y, n\}$.

Ex-ante, the probability for an agent i to be the high-cost type $\theta_i = \bar{\theta}$ is γ and conversely the probability to be the low-cost type $\theta_i = \underline{\theta}$ is $1 - \gamma$. These probabilities are independent of the other agents' types. Any initial type θ_i has the same probability δ of realising a positive shock $\varepsilon = y$ to their production cost. Conversely, with the same probability of $1 - \delta$ the agents realise no shock denoted $\varepsilon = n$. The probabilities γ and δ are identical between the agents and independent of the type and shock of both agents. We assume that the cost types are ordered such that $c(q, \underline{\theta}, \varepsilon) < c(q, \bar{\theta}, \varepsilon)$ ($\forall \varepsilon$ and q). Furthermore, without a shock ($\varepsilon = n$)

63

both types have a lower cost than with a shock, i.e. $c(q, \theta, n) < c(q, \theta, y)$ ($\forall \theta$ and q). Finally, we focus on the case where $c(q, \underline{\theta}, y) < c(q, \overline{\theta}, n)$, i.e. where the shock does not disturb the initial ordering of types. We can think of this as the initial type being more important in determining the cost type of an agent than the shock. This assumption is maintained in the motivating example as well as the general model.

A seller with type θ and shock ε makes quasilinear profits

$$\pi = p - c(q, \theta, \varepsilon) \tag{3.1}$$

of selling the good at price p and quality q. The buyer's preferences $U(p, q)$ over (p, q) are quasilinear of the form:

$$U(p, q) = V(q) - p. \tag{3.2}$$

In the event of no trade taking place, the buyer and the sellers all receive a payoff of 0. We assume that the cost $c(q, \theta, \varepsilon)$ is increasing in q ($c_q(q, \theta, \varepsilon) > 0$) at an increasing rate ($c_{qq}(q, \theta, \varepsilon) < 0$) for all θ, ε and $q > 0$. We additionally assume that there are always gains from trade to be made ($V(q) - c(q, \overline{\theta}, y) > 0$ for some q) and that $V(q)$ satisfies the Inada conditions to guarantee an interior solution, i.e. $\lim_{q \to 0} V'(q) = \infty$ and $\lim_{q \to \infty} V'(q) = 0$ are satisfied.

A general mechanism in this environment consists of an outcome function and messages from the agents to the principal before and after the shock.

The timing of events is as in chapter 2:

1. The principal commits to a mechanism.

2. The agents privately learn their cost type θ.

3. The agents choose a strategy denoting messages at each node of the game tree of the mechanism.

4. The mechanism may allocate the procurement of the good to an agent given the agents' messages.

5. If and only if an agent has been allocated the procurement contract, he privately realises the shock ε and thus knows his true cost parameter θ, ε. He may choose any additional message after the shock in response to this new private information.

6. The agent from step 5 produces the good for the principal given the outcome function and the agents' messages before and after the shock.

In the utilitarian social optimum, given the assumption of risk-neutrality, prices do not matter since they are a direct and costless transfer of utilities. Here, given the model assumptions, the contract should always be allocated to the seller with the lowest type θ and lowest shock ε. He will be asked to implement a quality such that the difference between the valuation of quality by the buyer and the cost of producing quality of the seller is maximised. For any cost type and shock, the socially optimal quality levels are given by the solutions to the set of equations:

$$V_q(q) = c_q(q, \theta, \varepsilon) \tag{3.3}$$

$(\forall \theta, \varepsilon)$. The qualities that solve these equations are referred to as the 'undistorted' or 'socially optimal' qualities.

The model described here is an abstraction of the complex procurement environment. There are N firms of 2 different possible types competing to sell a single two-dimensional good to the procurer. The uncertainty modelled is only one of two realisations for each of the types. This simplification is owed to the additional complexity introduced when discussing the optimal scoring auction. This model is an abstraction from the more complex continuous setting in chapter 2. The introduced stylisation will, however, allow us to extend the analysis conducted in that paper and can provide additional intuition to the continuous type results.

3.3   Motivating Example

To understand more clearly the model and results of our analysis, we present a motivating example below. This example lets $N = 2$ and defines specific functional forms satisfying our assumptions made above. In this way, we will be able to get numerical results for each of the findings presented in the paper. The example itself is based on the experience working with the National Health Service (NHS) procurement services in Scotland. The example is fictional. It is meant to illustrate critical issues that can arise if a procurement contract is complex, multidimensional, and dynamic. It is not meant to be an accurate representation of the procurement market for MRI scanners.

The NHS is comparing a centralised procurement process to letting each hospital manage its own purchasing of supplies. Over recent years, the NHS found that with many goods and services, it can achieve a better average price and lower administrative costs using the centralised procurement process compared to the decentralised scheme in place before. The NHS believes this gain in efficiency is primarily due to increased market power and reduced overhead costs. In this way, the centralised procurement process represents a less inefficient bureaucracy. Recent empirical evidence from Decarolis et al. (2018) for US federal procurement and Bucciol, Camboni, and Valbonesi (2020) suggest that a more competent procurement bureau can significantly reduce delays, cost-overruns, and the likelihood of costly renegotiation. A centralised scheme can be a positive factor moving in the direction of a more competent procurement bureau. Consequently, the NHS started to transition the purchasing of many basic goods and services, for example, the purchasing of sanitary products or cleaning and waste services, to a centralised procurement scheme. The mechanism the NHS uses for its procurement is a sealed bid scoring auction where firms submit a score made up of the combination of a price and a good that is rated on various quality aspects relevant to the particular context. This mechanism is very close to the first score auction as introduced in Che (1993), and for this reason, following Che's result, the NHS believes this

mechanism to be close to optimal for its purposes.

Due to the success of this process, the NHS is considering to switch more purchasing processes to a centralised procurement scheme. One of these decisions is the purchasing of MRI scanning equipment. The NHS knows that there are two large international companies supplying MRI scanners to its hospitals at the moment. Its preference over MRI scanners are two-dimensional, preferring a lower price (p) and a higher probability of accurate ailment detection and location per single scan ($q \in [0, 1]$). We refer to this probability as the *quality* of the MRI scanner. The NHS's preferences over p and q given by $U(p, q) = 2\sqrt{q} - p$. The NHS knows that the two available suppliers indexed $i = 1, 2$ make quasilinear profits of the form $\pi_i(p, q) = p - \frac{1}{2}(\theta_i + \varepsilon)q^2$ from supplying MRI scanners of quality q at price p. There are two factors contributing to a firm's overall cost of producing q. The firm's initial type θ_i and a cost shock ε. The NHS understands any of the firms may be of an efficient type with $\theta_i = 1$ with probability $1/3$ or an inefficient type with $\theta_i = 2$ with probability $2/3$. Furthermore, the reason the NHS is unsure about whether its existing procurement scheme will be efficient for the purpose of MRI scanners, is that this area is subject to frequent technology shocks ε. The NHS believes that in the time-frame after signing a procurement contract with a firm and before the supply of the scanners, there is a probability $\delta = 1/3$ of a positive cost shock to a computing component that is common to both firms' scanners. This shock is independent of which firm is selected for the procurement contract. The shock ε would raise the cost of producing any quality q for either type θ_i according to the cost function above by $\varepsilon = 0.5$.[1]

The NHS is interested in what the best procurement process for this context is. At the least, it would like to know whether the existing scoring auction scheme is optimal here.

---

[1]These numbers are chosen to ensure the qualities of the numerical example are monotone in the cost of producing quality throughout the paper. We ensure this for illustration purposes only and often do not need to assume such monotonicity. We do not need to assume this for the general model and for many reasonable model environments, the qualities will not be monotonic in this way.

Finally, given its previous success with the scoring auction, if the scoring auction is no longer optimal, the NHS would like to know how much is lost relative to the theoretic optimum if it were to continue to use the scoring auction.

3.4   The Optimal Mechanism

The principal is designing a mechanism to maximise her expected utility. A mechanism here defines to the agent for every realisation of the agent's strategy a probability of being allocated the contract x_i, transfers to the agent from the principal t_i, and a quality q to be implemented by the agent chosen for the second stage. To capture the idea of reagreeing the contract with a single agent, we are only interested in mechanisms in which the allocation choice is made before the realisation of the cost shock. The revelation principle applies in this context, as shown in chapter 2, lemma 2.3.1. We can thus restrict attention to direct mechanisms where the agents directly and truthfully reveal their initial type θ_i to the principal first, and the winning agent directly and truthfully reveals their shock ε_i to the principal in the second stage subject to truthtelling and participation constraints. The principal may choose to disclose to the winning agent information about the losing agents' types. However, given we are in a private cost setting, she does not gain anything by doing so.

3.4.1   The Optimisation Problem

There are two stages to the direct mechanism. The first stage comprises an initial transfer t_i^1 and a decision rule x_i for the allocation of the second stage in response to the revelation of the initial types. In the subsequent stage there may be further transfers t_i^2 and the qualities q are chosen in response to the revelation of the cost shock. Let $\boldsymbol{\theta} = \{\theta_1, .., \theta_N\}$ be the vector of the agents' types. The principal's optimisation programme is to choose $\{x_i(\boldsymbol{\theta}), t_i^1(\boldsymbol{\theta}), t_i^2(\boldsymbol{\theta}, \varepsilon), q(\boldsymbol{\theta}, \varepsilon)\}$ for all $i, \boldsymbol{\theta}$ and ε to maximise her expected utility subject to

incentive and participation constraints. The agent's total transfers from the mechanism are denoted $t_i(\boldsymbol{\theta}, \varepsilon) = t_i^1(\boldsymbol{\theta}) + t_i^2(\boldsymbol{\theta}, \varepsilon) x_i(\boldsymbol{\theta})$. The principal thus solves:

$$\max_{\{x,t,q\}} \sum_{i=1}^{N} E_{\boldsymbol{\theta}, \varepsilon} \Big[V(q(\boldsymbol{\theta}, \varepsilon)) x_i(\boldsymbol{\theta}) - t_i(\boldsymbol{\theta}, \varepsilon) \Big] \tag{3.4}$$

Subject to the first stage participation constraints:

$$E_{\boldsymbol{\theta}_{-i}, \varepsilon} \Big[t_i(\theta_i, \boldsymbol{\theta}_{-i}, \varepsilon) - c(q(\theta_i, \boldsymbol{\theta}_{-i}, \varepsilon), \theta_i, \varepsilon) x_i(\theta_i, \boldsymbol{\theta}_{-i}) \Big] \geq 0 \tag{3.5}$$

$(\forall i, \theta_i)$. Subject to the first stage incentive constraints:

$$\theta_i \in \operatorname*{argmax}_{\widetilde{\theta}_i} E_{\boldsymbol{\theta}_{-i}, \varepsilon} \Big[v_i(\widetilde{\theta}_i, \boldsymbol{\theta}_{-i}, \varepsilon) x_i(\widetilde{\theta}_i, \boldsymbol{\theta}_{-i}) + t_i^1(\widetilde{\theta}_i, \boldsymbol{\theta}_{-i}) \Big] \tag{3.6}$$

$(\forall i, \theta_i)$, where:

$$v_i(\widetilde{\theta}_i, \theta_i, \varepsilon) = \max_{\widetilde{\varepsilon}} E_{\boldsymbol{\theta}_{-i}} \Big[t_i^2(\widetilde{\theta}_i, \boldsymbol{\theta}_{-i}, \widetilde{\varepsilon}) - c(q(\widetilde{\theta}_i, \boldsymbol{\theta}_{-i}, \widetilde{\varepsilon}), \theta_i, \varepsilon) \Big| I_i \Big] \tag{3.7}$$

$(\forall i, \theta_i, \varepsilon, \widetilde{\theta}_i)$. Here I_i denotes the information the winning agent i may learn about the other agent's type conditional on winning the second stage[2]. Finally, the optimisation is subject to second stage incentive constraints:

$$\varepsilon \in \operatorname*{argmax}_{\widetilde{\varepsilon}} E_{\boldsymbol{\theta}_{-i}} \Big[t_i^2(\theta_i, \boldsymbol{\theta}_{-i}, \widetilde{\varepsilon}) - c(q(\theta_i, \boldsymbol{\theta}_{-i}, \widetilde{\varepsilon}), \theta_i, \varepsilon) \Big| I_i \Big] \tag{3.8}$$

$(\forall i, \varepsilon)$. Constraint (3.5) dictates that an agent i of either type θ_i should not make a loss from the mechanism in expectation over the shock. In other words, the agent has an outside option of zero. Note that we are not allowing the agent and principal to re-agree whether the good will be produced at all but only allow a change in price and quality. As such, the agent may thus incur a loss ex-post[3]. The constraints in equation (3.6) and (3.7) say that the agent of any type θ_i should see it profitable to reveal his true type in the first stage *even*

---

[2]This will turn out to be unimportant in equilibrium but is included here for completeness.

[3]Due to the risk-neutrality of the agents and the principal, voluntary participation in the second stage can be achieved with an increase in transfers after the shock compensated by a decrease in transfers before the shock. This is thus not restrictive.

if they are allowed to double deviate (i.e. lie about θ_i planning to also lie about ε). Given this first stage truthtelling constraint, the constraint (3.8) about second stage truthtelling states that the agent of type θ_i, if allocated the second stage, has to want to reveal his shock ε truthfully if there was first stage truthtelling about θ_i. Jointly these constraints ensure truthtelling on the equilibrium path but do not enforce truthtelling off the path.

3.4.2 The Optimal Mechanism

We impose the following assumption on the cost functions and model parameters:

Assumption 5. Monotonicity 1: *The optimal quality policies q^*, which form part of the solution to the mechanism design problem, are non-increasing in ε and θ separately. I.e. $q^*(\underline{\theta}, \varepsilon) \geq q^*(\overline{\theta}, \varepsilon)$ ($\forall \varepsilon$) and $q^*(\theta, n) \geq q^*(\theta, y)$ ($\forall \theta$).*

This assumption guarantees that the quality policies of the buyer optimal mechanism are non-increasing in each of the two parameters determining the cost of the agent. This assumption implies some restriction on the functional form of the cost functions and utility functions of the agents and the principal and is sufficiently strong for the purpose of simplifying the analytical solution to the mechanism design programme. An assumption of this type is common in the literature and in this form is not particularly restrictive. In section (3.6) we will use a stronger monotonicity condition given by assumption 3.6.1.

Given the symmetry between the agents we may simplify the solution programme by dropping the i subscript. Further, let $x(\theta)$ and $t^1(\theta)$ be shorthand for $E_{\boldsymbol{\theta}_{-i}}[x(\theta_i, \boldsymbol{\theta}_{-i})]$ and $E_{\boldsymbol{\theta}_{-i}}[t^1(\theta_i, \boldsymbol{\theta}_{-i})]$. The optimisation problem can be further simplified by realising that an agent's optimal quality policies will not depend on the other agents' types in the optimal mechanism.[4] Finally, we can make the second stage transfers independent of the other agents' types. Due to the assumed independence and risk-neutrality, this is without loss of

---

[4]This is ensured by the assumptions on the principal's utility function and the agents' cost functions. See the appendix A.2 in chapter 2 for a more thorough exposition.

generality. Thus $t^2(\theta, \varepsilon)$ and $q(\theta, \varepsilon)$ will denote the second stage transfers and quality for an agent who announces type θ in the first stage and shock ε in the second. We can then write the representative agent's participation constraints from equation (3.5) for the unproductive type $\theta = \bar{\theta}$ as:

$$t^1(\bar{\theta}) + E_\varepsilon\left[t^2(\bar{\theta}, \varepsilon) - c(q(\bar{\theta}, \varepsilon), \bar{\theta}, \varepsilon)\right]x(\bar{\theta}) \geq 0. \tag{3.9}$$

The first stage IC from equation (3.6) for type $\theta = \underline{\theta}$ to ensure he does not wish to pretend to be type $\theta = \bar{\theta}$ may be expressed as:

$$t^1(\underline{\theta}) + E_\varepsilon\left[v(\underline{\theta}, \underline{\theta}, \varepsilon)x(\underline{\theta})\right] \geq t^1(\bar{\theta}) + E_\varepsilon\left[v(\bar{\theta}, \underline{\theta}, \varepsilon)x(\bar{\theta})\right]. \tag{3.10}$$

As a reminder, $v(\widehat{\theta}, \theta, \varepsilon)$ denotes the value of an agent of type θ and shock ε reaching the second stage after announcing $\widehat{\theta}$ in the first stage. It is straightforward to show that the relevant first stage incentive and participation constraints for the optimal direct mechanism are the ones in equations (3.9) and (3.10).[5] The second stage IC constraints from equation (3.8) for an agent without a shock $\varepsilon = n$ pretending to have received one can be written as:

$$t^2(\theta, n) - c(q(\theta, n), \theta, n) \geq t^2(\theta, y) - c(q(\theta, y), \theta, n) \tag{3.11}$$

$(\forall \theta)$. If the optimal quality policies are decreasing in the cost shock for all initial types θ, then the above will be the relevant second stage incentive constraint. Letting

$$t^2(\theta, y) - c(q(\theta, y), \theta, y) = 0 \tag{3.12}$$

for all θ is without loss of generality as by the risk-neutrality of the agents we may move any transfer from the first to the second stage and vice versa without affecting the optimisation

---

[5]This is a standard result. The least productive type will end up with their reservation utility (normalised to zero) in equilibrium, and we only need to ensure he does not get less (and thus participates). However, since the more productive type mimicking the less productive type end up with larger profits because of their lower costs, we need to ensure the more productive types do not want to mimic the less productive type (incentive compatibility).

programme as long as we are still ensuring the constraints. It follows from equation (3.11) that:

$$t^2(\theta, n) \geq c(q(\theta, n), \theta, n) + c(q(\theta, y), \theta, y) - c(q(\theta, y), \theta, n) \tag{3.13}$$

($\forall \theta$). Since the incentive constraint is satisfied at the lowest transfers when it is binding, it will be optimal for the principal to set:

$$t^2(\theta, n) = c(q(\theta, n), \theta, n) + c(q(\theta, y), \theta, y) - c(q(\theta, y), \theta, n) \tag{3.14}$$

($\forall \theta$). From equation (3.7) we know that $v(\theta, \theta, \varepsilon)$ is given by:

$$v(\theta, \theta, \varepsilon) = \max_{\widetilde{\varepsilon}} \left\{ t^2(\theta, \widetilde{\varepsilon}) - c(q(\theta, \widetilde{\varepsilon}), \theta, \varepsilon) \right\} \tag{3.15}$$

($\forall \theta, \varepsilon$). Using the second stage truthtelling constraint, we can write the value function v when telling the truth in the first stage such that $\widehat{\theta} = \theta$ as:

$$v(\theta, \theta, \varepsilon) = t^2(\theta, \varepsilon) - c(q(\theta, \varepsilon), \theta, \varepsilon) \tag{3.16}$$

($\forall \theta, \varepsilon$). From equation (3.12) we get:

$$v(\theta, \theta, y) = 0 \tag{3.17}$$

($\forall \theta$). And using (3.14) we get:

$$v(\theta, \theta, n) = c(q(\theta, n), \theta, n) + c(q(\theta, y), \theta, y) - c(q(\theta, y), \theta, n) - c(q(\theta, n), \theta, n) \tag{3.18}$$

$$= c(q(\theta, y), \theta, y) - c(q(\theta, y), \theta, n) \tag{3.19}$$

($\forall \theta$). Next, we want to use the first stage incentive constraints to learn more about the first stage transfers. Going back to and rewriting the first stage incentive constraint for an agent of type $\theta = \underline{\theta}$ pretending to be of type $\theta = \overline{\theta}$ we get:

$$t^1(\underline{\theta}) + E_\varepsilon \left[v(\underline{\theta}, \underline{\theta}, \varepsilon) x(\underline{\theta}) \right] \geq t^1(\overline{\theta}) + E_\varepsilon \left[v(\overline{\theta}, \underline{\theta}, \varepsilon) x(\overline{\theta}) \right]. \tag{3.20}$$

We next want to describe the agents' value functions when misreporting in the first stage. A sufficient condition for an agent of type $\theta = \underline{\theta}$, $\varepsilon = y$ to prefer reporting $\varepsilon = n$ in the second stage if they misreported $\theta = \overline{\theta}$ in the first stage is given by:

$$t^2(\overline{\theta}, n) - c(q(\overline{\theta}, n), \underline{\theta}, y) > t^2(\overline{\theta}, y) - c(q(\overline{\theta}, y), \underline{\theta}, y). \tag{3.21}$$

We can substitute for the second stage transfers to get:

$$c(q(\overline{\theta}, n), \overline{\theta}, n) - c(q(\overline{\theta}, n), \underline{\theta}, y) > c(q(\overline{\theta}, y), \overline{\theta}, n) - c(q(\overline{\theta}, y), \underline{\theta}, y). \tag{3.22}$$

Since $c(q, \underline{\theta}, y) < c(q, \overline{\theta}, n)$ for any q, and since $q(\overline{\theta}, n) > q(\overline{\theta}, y)$ from the assumed monotonicity in ε, by the convexity of the cost function the inequality holds. Similarly, for an agent of type $\theta = \underline{\theta}, \varepsilon = n$ reporting $\theta = \overline{\theta}$ in the first stage we need:

$$c(q(\overline{\theta}, n), \overline{\theta}, n) - c(q(\overline{\theta}, n), \underline{\theta}, n) > c(q(\overline{\theta}, y), \overline{\theta}, n) - c(q(\overline{\theta}, y), \underline{\theta}, n) \tag{3.23}$$

to hold for him to report $\varepsilon = n$ in the second stage. Using the same logic as before we can show that it does. Therefore, any agent of type $\theta = \underline{\theta}$ misreporting to $\theta = \overline{\theta}$ prefers to report $\varepsilon = n$ in the second stage no matter whether they realise $\varepsilon = n$ or y. Thus:

$$v(\overline{\theta}, \underline{\theta}, n) = t^2(\overline{\theta}, n) - c(q(\overline{\theta}, n), \underline{\theta}, n) \tag{3.24}$$
$$= c(q(\overline{\theta}, n), \overline{\theta}, n) + c(q(\overline{\theta}, y), \overline{\theta}, y) - c(q(\overline{\theta}, y), \overline{\theta}, n) - c(q(\overline{\theta}, n), \underline{\theta}, n). \tag{3.25}$$

For $\varepsilon = y$ we get:

$$v(\overline{\theta}, \underline{\theta}, y) = t^2(\overline{\theta}, n) - c(q(\overline{\theta}, n), \underline{\theta}, y) \tag{3.26}$$
$$= c(q(\overline{\theta}, n), \overline{\theta}, n) + c(q(\overline{\theta}, y), \overline{\theta}, y) - c(q(\overline{\theta}, y), \overline{\theta}, n) - c(q(\overline{\theta}, n), \underline{\theta}, y) \tag{3.27}$$

We now need to obtain an expression for the first stage incentive constraints with only one transfer. To this end, we can see that cheapest way for the principal to satisfy the participation constraint of an agent of the least productive type $\theta = \overline{\theta}$ in equation (3.9) is at binding. Thus using (3.9) and (3.16) we get:

$$t^1(\overline{\theta}) = -E_\varepsilon \left[v(\overline{\theta}, \overline{\theta}, \varepsilon) \right] x(\overline{\theta}). \tag{3.28}$$

The productive type $\theta = \underline{\theta}$ will be able to extract informational rent from the mechanism in expectation. Thus the only participation constraint we need to satisfy is the above. For an agent of type $\theta = \underline{\theta}$, we need to ensure he does not wish to report being type $\theta = \overline{\theta}$. Substituting for $t^1(\overline{\theta})$ into the incentive constraint of equation (3.20) for $\theta = \underline{\theta}$ we get:

$$t^1(\underline{\theta}) = -E_\varepsilon\Big[v(\underline{\theta},\underline{\theta},\varepsilon)x(\underline{\theta})\Big] + E_\varepsilon\Big[v(\overline{\theta},\underline{\theta},\varepsilon) - v(\overline{\theta},\overline{\theta},\varepsilon)\Big]x(\overline{\theta}). \tag{3.29}$$

Here we used again the fact that the cheapest way for the principal to ensure incentive compatibility is at binding. By the convexity of the cost function, the constraints in equation (3.28) and (3.29) are sufficient to ensure participation and incentive compatibility if the expected quality policies (over the shock ε) are monotonic in the cost of producing q. We can further simplify these constraints using equation (3.17); however to save on notation, we proceed with (3.29) for now. Substituting for first and second stage transfers from equations (3.12), (3.14), (3.28), and (3.29) into the optimisation programme we get:

$$\max_{\{x,q\}} N \cdot E_\theta\Bigg[\Big((1-\delta)\Big(V(q(\theta,n)) - c(q(\theta,n),\theta,n) - c(q(\theta,y),\theta,y) + c(q(\theta,y),\theta,n)\Big)$$
$$+\delta\Big(V(q(\theta,y)) - c(q(\theta,y),\theta,y)\Big)\Big)x(\theta) + E_\varepsilon\Big[v(\theta,\theta,\varepsilon)x(\theta)\Big]\Bigg] \tag{3.30}$$
$$-(1-\gamma)E_\varepsilon\Big[v(\overline{\theta},\underline{\theta},\varepsilon) - v(\overline{\theta},\overline{\theta},\varepsilon)\Big]x(\overline{\theta}).$$

Note that $E_\varepsilon\Big[v(\theta,\theta,\varepsilon)x(\theta)\Big] = (1-\delta)\Big(c(q(\theta,y),\theta,y) - c(q(\theta,y),\theta,n)\Big)$ and that we can ignore any multiplicative scalar. This equation thus simplifies to:

$$\max_{\{x,q\}} E_{\theta\varepsilon}\Bigg[\Big(V(q(\theta,\varepsilon)) - c(q(\theta,\varepsilon),\theta,\varepsilon)\Big)x(\theta)\Bigg] - (1-\gamma)E_\varepsilon\Big[v(\overline{\theta},\underline{\theta},\varepsilon) - v(\overline{\theta},\overline{\theta},\varepsilon)\Big]x(\overline{\theta}). \tag{3.31}$$

The principal thus maximises the difference between the total expected social surplus from the procurement project:

$$E_{\theta\varepsilon}\Bigg[\Big(V(q(\theta,\varepsilon)) - c(q(\theta,\varepsilon),\theta,\varepsilon)\Big)x(\theta)\Bigg] \tag{3.32}$$

and the total expected informational rent extracted by the agent:

$$(1-\gamma)E_\varepsilon\Big[v(\overline{\theta},\underline{\theta},\varepsilon) - v(\overline{\theta},\overline{\theta},\varepsilon)\Big]x(\overline{\theta}) \tag{3.33}$$

for each agent. Taking the first order conditions of this optimisation programme with respect to the qualities q and collecting terms we find:

$$V'(q_{\underline{\theta}n}^*) = c_q(q_{\underline{\theta}n}^*, \underline{\theta}, n) \tag{3.34}$$

$$V'(q_{\underline{\theta}y}^*) = c_q(q_{\underline{\theta}y}^*, \underline{\theta}, y) \tag{3.35}$$

$$V'(q_{\overline{\theta}n}^*) = c_q(q_{\overline{\theta}n}^*, \overline{\theta}, n) + \frac{1-\gamma}{\gamma}\left(c_q(q_{\overline{\theta}n}^*, \overline{\theta}, n) - c_q(q_{\overline{\theta}n}^*, \underline{\theta}, n) \right.$$

$$\left. + \frac{\delta}{1-\delta}\left(c_q(q_{\overline{\theta}n}^*, \overline{\theta}, n) - c_q(q_{\overline{\theta}n}^*, \underline{\theta}, y) \right) \right) \tag{3.36}$$

$$V'(q_{\overline{\theta}y}^*) = c_q(q_{\overline{\theta}y}^*, \overline{\theta}, y) + \frac{1-\gamma}{\gamma}\left(c_q(q_{\overline{\theta}y}^*, \overline{\theta}, y) - c_q(q_{\overline{\theta}y}^*, \overline{\theta}, n) \right) \tag{3.37}$$

as necessary conditions for an interior solution to the optimal mechanism design programme under the assumptions made (for any given optimal allocation rule).

Proposition 3.4.1. *For any given optimal transfers and allocation rule, the qualities $q_{\theta\varepsilon}^*$ given in equation (3.34) to (3.37) are the quality policies of the optimal mechanism.*

Proof. We have shown above that these conditions are part of any interior solution given the model assumptions above. By the assumed Inada conditions on $V(q)$, an interior solution is guaranteed, and the conditions are thus both necessary and sufficient. $\qquad\square$

As we can see, in this simple dynamic model, the 'squeeze' of the quality policies for the type $\theta = \overline{\theta}, \varepsilon = n$, i.e. the reduction of his buyer-optimal qualities against the socially optimal ones, features an additional term we cannot find in comparable analyses. For example, a key result of the seminal dynamic mechanism design analysis by Eső and Szentes (2007) is that the agents' second stage private information does not affect the agents' informational rent. Our paper mirrors this, since the agent gives up all the additional informational rent he obtains from receiving no shock in the second stage through first stage transfers. Thus the two value functions $v(\theta, \theta, n)$ exactly cancel out in the simplification of the optimisation programme in equation (3.31). However, crucially this informational rent also becomes

relevant to the overall rent of an agent as it affects the profitability of a deviation to some other type $\widetilde{\theta}$ in the first stage of the direct mechanism. In particular, the productive agent deviating to the unproductive type in their first stage report prefers to report the low shock no matter what shock they realise as shown in equation (3.22) and (3.23). As a consequence, holding all else equal, in order to decrease the productive agent's rent, the principal should punish the unproductive type with the shock less than without the shock. As such, even though like previous analyses on related topics such as Eső and Szentes (2007) it is the first stage private information that primarily determines the agent's rent, the private information of the second stage is not irrelevant since it determines the value of first stage deviations. Since this expected value of a first stage deviation is nothing but the expected informational rent an agent can obtain from the mechanism, the second stage private information directly matters to the informational rent an agent can extract.

3.4.3 Numerical Example

In the case of the NHS, where $N = 2$ and the functions and probabilities are given as in section 3.3, the socially optimal qualities $q^o_{\theta\varepsilon}$ for agents of type θ with shock ε that maximise social welfare defined by equation (3.3) can be found as:

$$q^o_{\underline{\theta}n} = 100\%, \ q^o_{\underline{\theta}y} = 76\%, \ q^o_{\bar{\theta}n} = 63\%, \ \text{and} \ q^o_{\bar{\theta}y} = 54\%. \tag{3.38}$$

Here we expressed the qualities as percentage probabilities of ailment discovery and location. Note that $U(p, 1) = 2 - p$ and $U(p, 0) = -p$. Thus the NHS values the difference between a 100% and a 0% quality of the scanner at 2 units of p. We can apply the functional forms of the numerical example to the solution of the generalised direct mechanism optimisation problem to find the buyer optimal qualities $q^*_{\theta\varepsilon}$ for agents of type θ with shock ε:

$$q^*_{\underline{\theta}n} = 100\%, \ q^*_{\underline{\theta}y} = 76\%, \ q^*_{\bar{\theta}n} = 53\%, \ \text{and} \ q^*_{\bar{\theta}y} = 51\%. \tag{3.39}$$

We can see that in the optimal mechanism, the principal does not wish to distort the qualities of the most productive firm with or without the shock. From equations (3.34) and (3.35),

this is a general result for the $N \times 2$ type-space. This also extends to a continuous type-space with a continuous shock, as seen in chapter 2. Furthermore, the absolute and relative decrease in the optimal mechanism qualities away from the socially optimal ones is larger for the less productive types without the shock than with the shock. This can be attributed to the additional squeeze present in equation (3.36) but not (3.37). Intuitively, the principal can decrease the informational rent paid to the agents by making a deviation to the less productive type worse. As we learnt, given monotonicity in q, the more productive type will always pretend not to have realised the cost shock in the second stage if they pretend to be less productive in the first stage. As such, in order to make a deviation for these types worse, the optimal mechanism squeezes the quality $q^*_{\underline{\theta}n}$ more than $q^*_{\underline{\theta}y}$.

3.5   Scoring Auctions with Renegotiation (SAWR)

Now that we understand the properties of the theoretically optimal procurement mechanism for the NHS, we want to know whether the current system of a scoring auction can implement it. The NHS is planning to run a sealed bid second-score auction and respond to potential cost shocks by renegotiating the contract. We restrict attention to second-score auctions since by the result of Che (1993) and chapter 2, scoring auctions are revenue equivalent in first- and second-score without and with renegotiation. Since in this chapter we are facing a discrete type space, restricting attention to second-score auction allows us to focus on pure strategies and we do not need to consider mixed strategies for the equilibrium. Since the uncertainty in our example is on the supply side, the NHS is planning to allow the supplier to unilaterally make an offer to the NHS, which it may choose to accept or reject. In order to not be seen to favour any firm in the procurement process unfairly, the decision rule for accepting or rejecting a contract amendment is that the new offer must achieve the same score of the previous contract to be accepted. A description of the implementation game for the general $N \times 2$ setup is given below.

3.5.1 The SAWR Implementation Game for Discrete Type-Spaces

The SAWR implementation has the following timeline of events:

1. The N agents realise their types $\boldsymbol{\theta} = \{\theta_1, \cdots, \theta_N\}$.

2. The principal announces an auction with scoring rule $S(p, q)$. Given the discrete type-space it makes sense to restrict attention to second-score auctions.

3. The agents simultaneously bid scores $\boldsymbol{S} = \{S_1, \cdots, S_N\}$.

4. The agent with the highest score bid wins the auction at S^1 with the second-highest score bid being S^2. If $S^1 = S^2$, the principal uses a tiebreaker to randomly select one of the agents bidding this score with equal probability.

5. Before any implementation of the qualities, the winning agent realises the shock ε.

6. The winning agent may choose a price p and quality q such that $S(p, q) \geq S^2$.

7. The p and q are implemented.

This model is adapted from the original setup of the scoring auction in Che (1993). The restriction to second-score auctions in the discrete type environment simplifies the analysis as we do not need to consider mixed strategies. Point 3 simplifies the bidding to just a score S. In our model, we assume that no price and quality are implemented before the shock is realised in point 5. Given this stylisation of events after the allocation decision is made, the initial bid only matters in deciding the winner and setting the score to be satisfied in the renegotiation stage. As such, the composition of the bid score in terms of p and q does not matter. Point 4 restricts attention to efficient auctions in which the good is always allocated. In the optimal mechanism, this assumption is not restrictive. Point 6 explains the basic second-score auction. We will see in section 3.6.2 that the principal can do better by setting a reserve score system. However, we do not consider this at this point. Point 7

just states that the payoffs for the principal and the agents are realised at the end of the mechanism, and the good is not supplied before the realisation of the shock.

We want to check whether the optimal mechanism of section 3.4 can be implemented using the SAWR game described here for the general model as well as our motivating example.

3.5.2 Equilibrium of second-score SAWR in discrete Type-Spaces

We restrict attention to quasilinear scoring rules that are linear in price in the same way as Che (1993) and in chapter 2. Furthermore, in discrete type-spaces it makes sense to focus on scoring rules that take the form of a step function:

Assumption 6. $S(p, q) = s(q) - p$, where $s(q)$ is an increasing step function, i.e. a linear combination of finite number of indicator functions for intervals such that $s(q) \geq s(q')$ if $q > q'$.

These type of scoring rules very closely resemble scoring rules used in procurement environments in practice. Here the score is usually linearly decreasing in the price, and quality aspects are awarded an increase in score for reaching certain levels. These scoring rules further ensure that agents do not deviate to qualities different from the optimal qualities suggested by the optimal mechanism solution. When there is only one quality dimension setting $s(0) = 0$ works as a minimum standard since $V(0) = 0$.

To describe the equilibrium of SAWR it will be helpful to introduce two concepts. First, an isoscore curve is a combination of p and q that achieve the same score S. Second, an isoprofit curve is a combination of p and q that achieve the same profit $\pi(\theta, \varepsilon)$ for an agent of cost type θ and shock ε. In equilibrium of the second stage of SAWR an agent of type θ, ε maximises their profits $\pi = p - c(q, \theta, \varepsilon)$ subject to satisfying the same score restriction $S(p, q) = S_2$. Substituting in this simplifies to the supplier choosing:

$$q(\theta, \varepsilon) = \operatorname*{argmax}_{q} s(q) - c(q, \theta, \varepsilon) \tag{3.40}$$
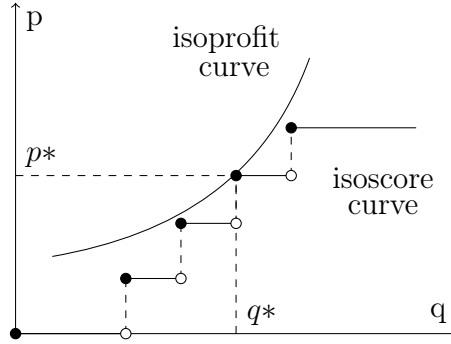
Figure 3.1 Isoscore and Isoprofit Curves with Discrete Types

subject to $S(p, q) = S_2$. Graphically, this is represented in figure 3.1. It is important to distinguish the quality defined by equation (3.40) from the 'pseudotype' as defined by Asker and Cantillon (2008) which looks deceptively similar. The former is the optimal quality that a seller in the second stage will select once they realised the shock whereas the latter is the quasi social surplus generated by a seller for a given scoring rule once they have chosen the quality with equation (3.40). By the assumptions on the buyer's cost function, the isoprofit curves will be strictly convex in the (p, q) space and given assumption 6, the isoscore curves are increasing step functions as shown here. By the assumption of quasilinearity for the cost functions and scoring rules, both the map of the isoprofit curves and the map of the isoscore curves are parallel shifts up/down the p axis. The best way to achieve any score $\overline{S}$ given the step function isoscore curves is to select the quality and price q^*, p^* that are on the highest isoprofit curve. Note, this will necessarily be at one of the black dots depicted in figure 3.1. Furthermore, the optimal quality q^* will be the same for a given cost function no matter the score S required to be matched by the price and quality. That is, the score to be matched only determines the price paid to the agent but not the quality the agent will choose to implement.

In equilibrium of second-scoreauctions, the agent will thus choose the price $p(\theta, \varepsilon)$ in the second stage as:

$$p(\theta, \varepsilon) = s(q(\theta, \varepsilon)) - \overline{S}. \tag{3.41}$$

80

Here $\overline{S}$ is the score required to be matched in the second-score auction and $q(\theta, \varepsilon)$ is the quality from equation (3.40). In the discrete type-space the equilibrium of SAWR is described by the following proposition:

Proposition 3.5.1. *A symmetric equilibrium of an efficient second-score SAWR mechanism[6] with discrete types is one in which each winning firm of type θ, ε chooses:*

$$q(\theta, \varepsilon) = \underset{q}{\operatorname{argmax}}\, s(q) - c(q, \theta, \varepsilon) \tag{3.42}$$

and:

$$p(\theta, \varepsilon) = s(q(\theta, \varepsilon)) - \overline{S} \tag{3.43}$$

where $\overline{S}$ is the score required to be implemented by the agent in the second stage. The equilibrium price choice is, exactly as in Che (1993) and Asker and Cantillon (2008) thus determined after the quality. There is a unique dominant strategy equilibrium of the first stage where each agent of type θ bids:

$$\widehat{S}(\theta) = E_\varepsilon \Big[s(q(\theta, \varepsilon)) - c(q(\theta, \varepsilon), \theta, \varepsilon) \Big]. \tag{3.44}$$

Proof. The second stage strategies are equilibrium strategies by the argument outlined above the proposition. To see the first stage strategies first note that neither type has any incentive to bid a different score. This is a straightforward extension of Vickrey (1961) and Che (1993). For instance, a type θ has no incentive to bid higher as although he wins the auction more often, those additional times where he wins, he makes a loss in expectation over ε. $\qquad\square$

3.5.3 The Implementation Problem

We can show that no scoring rule can implement the optimal mechanism by considering conditions that such a scoring rule would have to meet and illustrating that satisfying these properties would lead to higher expected transfers than the optimal mechanism.

---

[6]I.e. one where the contract is always allocated.

Proposition 3.5.2. *When $c(q, \underline{\theta}, y) < c(q, \overline{\theta}, n)$, no scoring rule $S(p, q)$ satisfying assumption 6 exists with which the SAWR implementation game detailed above implements the optimal mechanism for all possible model environments.*

Proof. This proof is for when the quality policies are monotonic in the cost of producing quality and not just in θ and ε separately. We know that such cases exist, since in the numerical example, the qualities are monotonic in the cost of producing q. To show that this proposition holds, we will begin by assuming that it does not and that some scoring rule $S(p, q) = s(q) - p$ exists that implements the optimal mechanism. For the principal to incentivise agents to produce qualities at the levels from the optimal mechanism in the equations (3.34) to (3.37) she can use scoring rules whose isoscore curves are step functions. Here, if the steps occur at the qualities of the optimal mechanism, the agents will always procure one of these qualities in equilibrium. To ensure the agents procure the qualities designed for them, the $s(q)$ functions will have to satisfy incentive constraints. For an agent of type θ, ε to produce quality $q_{\theta\varepsilon}^*$ in equilibrium of SAWR it has to hold that:

$$s(q(\theta, \varepsilon)) - c(q, \theta, \varepsilon) \geq s(\widetilde{q}) - c(\widetilde{q}, \theta, \varepsilon) \tag{3.45}$$

Note that in the order $\{(\underline{\theta}, n), (\underline{\theta}, y), (\overline{\theta}, n), (\overline{\theta}, y)\}$ the cost functions are monotonically increasing. If the quality policies of the optimal mechanism are monotonically decreasing, since the cost functions are convex, the principal can achieve full separation when the following set of equations are satisfied:

$$s(q_{\overline{\theta}y}^*) - c(q_{\overline{\theta}y}^*, \overline{\theta}, y) = \underline{S}, \tag{3.46}$$

$$s(q_{\overline{\theta}n}^*) - c(q_{\overline{\theta}n}^*, \overline{\theta}, n) = s(q_{\overline{\theta}y}^*) - c(q_{\overline{\theta}y}^*, \overline{\theta}, n), \tag{3.47}$$

$$s(q_{\underline{\theta}y}^*) - c(q_{\underline{\theta}y}^*, \underline{\theta}, y) = s(q_{\overline{\theta}n}^*) - c(q_{\overline{\theta}n}^*, \underline{\theta}, y), \tag{3.48}$$

$$\text{and } s(q_{\underline{\theta}n}^*) - c(q_{\underline{\theta}n}^*, \underline{\theta}, n) = s(q_{\underline{\theta}y}^*) - c(q_{\underline{\theta}y}^*, \underline{\theta}, n). \tag{3.49}$$

Here $\underline{S} \in \mathbb{R}^+$ is an arbitrary positive constant. These equations exactly identify scores $s(\cdot)$ for the qualities $q_{\underline{\theta}n}^*$ through to $q_{\overline{\theta}y}^*$ from equations (3.34) to (3.37). Raising or lowering the

score for all qualities by the same constant has no effect in equilibrium as the score of one quality only matters relative to the score of other qualities and has no intrinsic benefit to the principal or the agents and we can set $\underline{S}$ arbitrarily. For a given $\underline{S}$, the principal the principal wants to minimise the scores awarded to each of the qualities whilst maintaining incentive compatibility. To see why note that the buyer's expected utility is decreasing in the price and raising any score above the one determined by the above equations (weakly) increases the expected price paid to the agents. Equations (3.46) to (3.49) thus are necessary and sufficient conditions to fix the optimal scoring rule $S(p, q)$ under assumption 6.

To show that the SAWR implementation game does not implement the optimal mechanism, we consider the expected utility of the principal in the optimal mechanism as well as the SAWR implementation game. Given the equilibrium bidding strategy of the agents outlined above, the principal earns

$$E_{\boldsymbol{\theta}\varepsilon}[U(p,q)] = E_{\boldsymbol{\theta}\varepsilon}\left[\left(V(q^*_{\theta\varepsilon}) - c(q^*_{\theta\varepsilon}, \theta, \varepsilon)\right)\text{Prob}(\theta \in \min\{\boldsymbol{\theta}\})\right] - N \cdot E_{\theta\varepsilon}[\pi^{Sc}(\theta)] \quad (3.50)$$

from each agent. Here

$$E_{\theta\varepsilon}[\pi^{Sc}(\theta)] = (1 - \gamma)E_{\varepsilon}[\pi^{Sc}(\underline{\theta})] \quad (3.51)$$

$$= (1 - \gamma)\text{Prob}(\boldsymbol{\theta_{\text{-}i}} = \{\overline{\theta}\}^{N-1})E_{\varepsilon}\left[\widehat{S}(\underline{\theta}) - \widehat{S}(\overline{\theta})\right] \quad (3.52)$$

are the expected profits of an agent in SAWR before realising their type θ. $\widehat{S}(\theta)$ is the equilibrium score bid by an agent of type θ as defined in equation (3.44). Using the definition of $\widehat{S}(\theta)$ in equation (3.44) we get:

$$E_{\varepsilon}\left[\widehat{S}(\underline{\theta}) - \widehat{S}(\overline{\theta})\right] = (1 - \delta)\left(c(q^*_{\underline{\theta}y}, \underline{\theta}, y) - c(q^*_{\underline{\theta}y}, \underline{\theta}, n)\right) + \left(c(q^*_{\overline{\theta}n}, \overline{\theta}, n) - c(q^*_{\overline{\theta}n}, \underline{\theta}, y)\right)$$
$$+ \delta\left(c(q^*_{\overline{\theta}y}, \overline{\theta}, y) - c(q^*_{\overline{\theta}y}, \overline{\theta}, n)\right). \quad (3.53)$$

In the optimal mechanism, the principal earns

$$E_{\boldsymbol{\theta}\varepsilon}[U] = E_{\boldsymbol{\theta}\varepsilon}\left[\left(V(q^*_{\theta\varepsilon}) - c(q^*_{\theta\varepsilon}, \theta, \varepsilon)\right)\text{Prob}(\theta \in \min\{\boldsymbol{\theta}\})\right]$$
$$- N(1 - \gamma)E_{\varepsilon}\left[v(\overline{\theta}), \underline{\theta}, \varepsilon) - v(\overline{\theta}), \overline{\theta}, \varepsilon)\right]x(\overline{\theta}). \quad (3.54)$$

To compare the expected payoffs to the principal of the two mechanisms first note that the two expressions end with the identical social surplus but differ in how they allocate rent to the agents. We can write the rent term from the second line of equation (3.54) in terms of the cost functions of the agents. First, note that using equations (3.17), (3.19), (3.25), and (3.27) we get:

$$E_\varepsilon\left[v(\bar{\theta}),\underline{\theta},\varepsilon) - v(\bar{\theta}),\bar{\theta},\varepsilon)\right] = (1-\delta)\left(c(q^*_{\bar{\theta}n},\bar{\theta},n) - c(q^*_{\bar{\theta}n},\underline{\theta},n)\right) + \delta\left(c(q^*_{\bar{\theta}n},\bar{\theta},n) - c(q^*_{\bar{\theta}n},\underline{\theta},y)\right.$$
$$\left. + c(q^*_{\bar{\theta}y},\bar{\theta},y) - c(q^*_{\bar{\theta}y},\bar{\theta},n)\right) \tag{3.55}$$

$(\forall\theta)$. Thus the rent term of equation (3.54) becomes:

$$N(1-\gamma)\left(E_\varepsilon\left[c(q^*_{\bar{\theta}n},\bar{\theta},n) - c(q^*_{\bar{\theta}n},\underline{\theta},\varepsilon)\right] + \delta\left(c(q^*_{\bar{\theta}y},\bar{\theta},y) - c(q^*_{\bar{\theta}y},\bar{\theta},n)\right)\right)x(\bar{\theta}). \tag{3.56}$$

In equilibrium of the optimal mechanism the probability $x(\bar{\theta})$ of winning the allocation if you report type $\theta = \bar{\theta}$ is the probability of there being no $\theta = \underline{\theta}$ among the other agents' reports multiplied by the probability of winning the tiebreaker $(1/N)$:

$$x(\bar{\theta}) = \frac{1}{N}\cdot\text{Prob}(\boldsymbol{\theta_{-i}} = \{\bar{\theta}\}^{N-1}). \tag{3.57}$$

Thus to compare the expected utility for the principal from the optimal mechanism and SAWR, we can take the difference between equation (3.50) and (3.54). If the following inequality holds the rent paid to the agents is larger in SAWR than the optimal mechanism:

$$N(1-\gamma)\text{Prob}(\boldsymbol{\theta_{-i}} = \{\bar{\theta}\}^{N-1})\left((1-\delta)\left(c(q^*_{\bar{\theta}y},\underline{\theta},y) - c(q^*_{\bar{\theta}y},\underline{\theta},n)\right) + \left(c(q^*_{\bar{\theta}n},\bar{\theta},n)\right.\right.$$

$$\left.\left. -c(q^*_{\bar{\theta}n},\underline{\theta},y)\right) + \delta\left(c(q^*_{\bar{\theta}y},\bar{\theta},y) - c(q^*_{\bar{\theta}y},\bar{\theta},n)\right)\right) > N(1-\gamma)\frac{1}{N}\text{Prob}(\boldsymbol{\theta_{-i}} = \{\bar{\theta}\}^{N-1})$$

$$\cdot\left(E_\varepsilon\left[c(q^*_{\bar{\theta}n},\bar{\theta},n) - c(q^*_{\bar{\theta}n},\underline{\theta},\varepsilon)\right] + \delta\left(c(q^*_{\bar{\theta}y},\bar{\theta},y) - c(q^*_{\bar{\theta}y},\bar{\theta},n)\right)\right) \tag{3.58}$$

It is straightforward to show the inequality holds. Furthermore, the SAWR mechanism is worse than the optimal mechanism in two ways. First of all, since

$$\text{Prob}(\boldsymbol{\theta_{-i}} = \{\bar{\theta}\}^{N-1}) > \frac{1}{N}\text{Prob}(\boldsymbol{\theta_{-i}} = \{\bar{\theta}\}^{N-1}), \tag{3.59}$$

in SAWR, the agents are paid the full difference between the productive type and the unproductive type if they are the only productive type. In the optimal mechanism, this is reduced by the probability of having to win a tiebreaker which the agent only does $1/N$ of the time in the scenario where all other agents are of the high-cost type. Furthermore, we can compare the right-hand side of the two informational rents on each side of the equality in equation (3.58) to find:

$$(1-\delta)\Big(c(q^*_{\underline{\theta}y}, \underline{\theta}, y) - c(q^*_{\underline{\theta}y}, \underline{\theta}, n)\Big) + \Big(c(q^*_{\bar{\theta}n}, \bar{\theta}, n) - c(q^*_{\bar{\theta}n}, \underline{\theta}, y)\Big) + \delta\Big(c(q^*_{\bar{\theta}y}, \bar{\theta}, y)$$
$$-c(q^*_{\bar{\theta}y}, \bar{\theta}, n)\Big) > \Big(E_\varepsilon\Big[c(q^*_{\bar{\theta}n}, \bar{\theta}, n) - c(q^*_{\bar{\theta}n}, \underline{\theta}, \varepsilon)\Big] + \delta\Big(c(q^*_{\bar{\theta}y}, \bar{\theta}, y) - c(q^*_{\bar{\theta}y}, \bar{\theta}, n)\Big). \qquad (3.60)$$

Cancelling out terms and rewriting, this inequality reduces to:

$$c(q^*_{\underline{\theta}y}, \underline{\theta}, y) - c(q^*_{\underline{\theta}y}, \underline{\theta}, n) > c(q^*_{\bar{\theta}n}, \underline{\theta}, y) - c(q^*_{\bar{\theta}n}, \underline{\theta}, n). \qquad (3.61)$$

By the assumed monotonicity of q in the cost of producing quality, the convexity of the cost function, and since $c(q, \underline{\theta}, y) > c(q, \underline{\theta}, n)$ for all $q > 0$, this inequality holds. $\qquad \square$

Intuitively, in the renegotiation stage, the SAWR implementation game does not make optimal use of the relevant information about the winning agent's type θ, which he already revealed in the first stage through their bidding behaviour. As such, the agent can extract more informational rent than in the optimal mechanism.

The proof for proposition 3.5.2 shows that when trying to implement the optimal mechanism of section 3.4 through scoring auctions with renegotiation the agents are able to extract excessive rent through two channels. First, as an artefact of the discrete type-space, in the optimal mechanism, deviation to worse initial types is discouraged by the possibility of tie-breakers arising in the allocation rule. By decreasing a deviating agent's probability of winning the contract when misreporting to a higher type, in the optimal mechanism we can decrease the transfers to the agent to take this into account. Second-scoreSAWR does not decrease the transfers to the agent to take into account that they have to face more tiebreakers should they deviate to a worse type. As such, the agent can extract more rent

in SAWR than in the optimal mechanism. This first shortcoming of SAWR relative to the optimal mechanism can be made up by introducing a reserve score system which effectively reallocates the excessive rent from not having to face tiebreakers to the principal. This will is described in section 3.6.2.

Unlike the issue with not accounting for the agent suffering from tiebreakers in the optimal mechanism, the second shortcoming of SAWR relative to the optimal mechanism is harder to overcome. Intuitively, the quality policies of the optimal mechanism, when implemented through SAWR, lead to excessive rent because the implicit incentive constraints of SAWR are stronger than the constraints of the optimal direct mechanism. This problem is worse than it seems. As we will show in the numerical section below, in many model environments, SAWR with the OM qualities but without a reserve score system is considerably worse off than in the optimal mechanism. Furthermore, it only marginally outperforms a naïve scoring rule where $S(p,q) = U(p,q)$.

3.5.4 Numerical Example

Using the above analysis, we are able to quantify the loss for the NHS due to implementing an inefficient scoring auction with renegotiation rather than the optimal mechanism. The expected utility from the optimal mechanism for the NHS is given by

$$E[U] = 1.279, \text{Total Rent} = 0.030, \text{Social Benefit} - \text{Cost} = 1.309. \tag{3.62}$$

For SAWR without a reserve score and with the qualities of the optimal mechanism we get:

$$E[U] = 1.226, \text{Total Rent} = 0.083, \text{Social Benefit} - \text{Cost} = 1.309. \tag{3.63}$$

Finally, for SAWR with a naïve scoring rule and no reserve score system we get:

$$E[U] = 1.217, \text{Total Rent} = 0.098, \text{Social Benefit} - \text{Cost} = 1.315. \tag{3.64}$$

In this particular example, therefore, the NHS loses 4.14% of it's expected utility from using the inefficient SAWR mechanism relative to the optimal mechanism. For the naïve scoring

rule, in this example, this number is even higher at a 4.88% loss of efficiency. The former constitutes the equivalent of lowering the probability of both agents being a low type from 1/3 to slightly under 1/5, over a 40% decrease. The final thing to note here is that the increase in expected utility of the principal from using SAWR with the optimal mechanism qualities versus a naïve SAWR is only 0.002 units. However, the former is computationally significantly more demanding and comes with a loss of 0.05 units of social surplus. A factor of 20 times the gain in expected utility.

We can also use these numbers to show what proportion of the loss is due to the inefficient reduction of the agents' rent through not accounting for tiebreakers and how much is due to the q maximising the wrong maximand. Assuming that SAWR with the optimal mechanism q allocated rent at the same probability as the optimal mechanism, the expected utilities become:

$$E[U] = 1.268, \text{Total Rent} = 0.042, \text{Social Benefit} - \text{Cost} = 1.309. \tag{3.65}$$

SAWR only loses 0.89% efficiency vs the optimal mechanism in this case, implying that roughly 79% of the loss in efficiency is due to SAWR not accounting for tiebreakers correctly and 21% due to the wrong q.

Fortunately for SAWR, we can overcome its apparent main shortcoming, ignoring the possibility of tiebreakers in the optimal mechanism, by introducing a reserve score system. We do this in section 3.6.2. However, the problem of the qualities is not as easily overcome. Furthermore, if the qualities of the OM were designed to optimally disturb the social surplus given a different set of incentive constraints than those pertinent to SAWR, then the qualities of the OM are not the optimal qualities for SAWR. As such, SAWR with the optimal mechanism qualities is not the best possible SAWR. This raises another question. If we are still interested in finding the best scoring auction under the presence of renegotiation, and we argue that given its popularity in practice we should be since we can no longer use the qualities of the optimal mechanism for the design of the scoring auction, what methodology can we use for finding the optimal SAWR as the best among a class of suboptimal mech-

anisms? In section 3.6, we propose one such methodology and use it to identify the best SAWR.

3.6   Restricted Optimal Mechanism Design

Rather than optimising over scoring auctions directly, a relaxed programme can be solved to derive the optimal scoring rule for the SAWR implementation game indirectly. In particular, it is possible to get towards the optimal SAWR mechanism using a direct mechanism design problem where an additional restriction is imposed to reduce the mechanism space. In figure 3.2 an illustration is provided for reference.
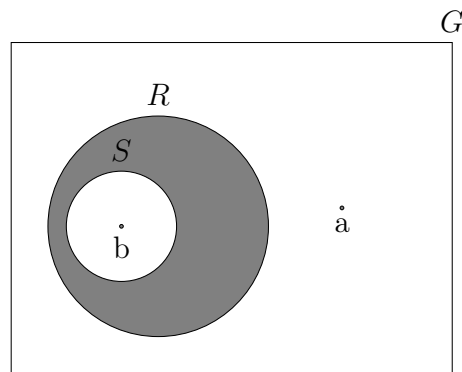
Figure 3.2 Illustration of Solution Methodology

Let the rectangle area G in the illustration be the set of social choice functions implementable in equilibrium by any general mechanism. By the revelation principle, any social choice function within this set can be implemented in truthtelling equilibrium with a direct mechanism of the form discussed. The buyer-optimal social choice function is marked by $a \in G$. The solution to the direct mechanism optimisation programme in section 3.4 gives the buyer optimal social choice function. The subset $S \subset G$ gives the (non-empty[7]) set of social choice functions that are implementable by SAWR implementation games with dif-

---

[7]To show that this is non-empty it is enough to convince oneself that SAWR with the 'naïve' scoring rule $S(p,q) = V(q) - p$ implements a socially optimal (but not buyer optimal) social choice function.

88

ferent scoring rules. From proposition 3.5.2, $a \notin S$. One approach to get towards $b \in S$, the social choice function implemented by the optimal SAWR mechanism, is to solve for the buyer-optimal scoring rule within SAWR directly. Such an optimisation problem is complicated: the optimal scoring rule as well as the optimal quality and price policies as functions of agents' bids would have to be identified simultaneously. As such, we do not pursue this approach in this paper but rather use an alternative approach suggested by figure 3.2 to get to b by solving a constrained mechanism design problem.

The aim is to identify a constraint on the mechanism design problem that allows searching for the optimal social choice function in the restricted subset $R \subset G$ rather than all of G. If within the restricted subset R, b is optimal, and if all of S is contained within R, it is possible to optimise over the restricted set to get towards b. A restriction on the direct mechanism design problem thus has to satisfy two key properties in order for this relaxed optimisation programme to work. First, the imposed restriction needs to be wide enough to not rule out any potential scoring rules of the SAWR game, i.e. the restriction needs to satisfy that (i) *Any SAWR equilibrium social choice function is implementable by a direct mechanism that satisfies the imposed restriction* ($S \subseteq R$). Second, the restriction also needs to ensure that R is narrow enough that b is optimal within R, i.e. it has to hold that (ii) *The social choice function of the optimal direct mechanism with the restriction is implementable by SAWR.* If we can find a restriction that satisfies these conditions, the optimal scoring rule (and thus the optimal SAWR mechanism) can be found as the scoring rule that implements the social choice function of the optimal mechanism from this restricted set.

In summary, rather than finding the optimal SAWR mechanism by optimising over scoring rules, the optimal SAWR is found by solving a direct mechanism programme to describe the properties any optimal SAWR mechanism would have to satisfy. The optimal SAWR is then identified by a scoring rule that implements these properties. This approach has the advantage of decomposing the optimisation problem into different components. First, the optimal qualities, transfers, and allocation decisions of the restricted optimal mechanism are

found. Second, a scoring auction is found that implements these.

A pertinent feature of the SAWR implementation game is that the renegotiation is set up in such a way that the principal does not make optimal use of the information revealed by the agent in the first stage. In terms of menus of options, in the second stage of SAWR, the agent's menu still includes all quality options q that were available in the first stage. If the principal wants to incentivise an agent to select a particular p, q bundle in the second stage, she needs to make sure that this agent prefers p, q over any other bundle that achieve at least the same score. The p, q have to be robust to deviations to other combinations. Not just for agents of the same type and a different shock, but also for agents of different types. In the optimal direct mechanism, however, the type θ is reported truthfully in the first stage and thus taken as given in the second stage. As such, the winning agent's second stage menu of options is limited to only those θ, ε bundles that are possible given their first stage report, i.e. θ, n and θ, y. It no longer includes the bundles of the other types $\widetilde{\theta} \neq \theta$. Given this feature of SAWR, we can guess a restriction on the direct mechanism design problem that may satisfy both (i) and (ii).

Guess Restriction 1. *Restrict the mechanism space to direct mechanisms which satisfy a stronger truthtelling constraint for the second stage:*

$$\theta_i, \varepsilon = \underset{\widetilde{\theta}_i, \widetilde{\varepsilon}}{\operatorname{argmax}} \left[t_i^2(\widetilde{\theta}_i, \widetilde{\varepsilon}) - c(q_i(\widetilde{\theta}_i, \widetilde{\varepsilon}), \theta_i, \varepsilon) \right] \tag{3.66}$$

($\forall i, \theta_i, \varepsilon$). I.e. direct mechanisms in which truthtelling is satisfied about both θ and ε in the second stage (not just about ε) and in which the second stage qualities and transfers may not depend on first stage reports.

The guess restriction can be shown to satisfy both features (i) and (ii) described above. First, for (i):

Lemma 3.6.1. *Any SAWR equilibrium social choice function is implementable by a restricted direct mechanism that enforces the guess restriction.*

The proof of lemma 3.6.1 involves adapting the proof of the revelation principle from chapter 2. However, rather than showing that any equilibrium outcome of any mechanism can be implemented by a direct mechanism, here we show that any equilibrium outcome of one particular mechanism can be implemented with a restricted direct mechanism. See appendix B.1 for a full derivation of this result.

To show that (ii) is satisfied by the guess restriction, we first derive the optimal *restricted* direct mechanism and second show that its outcome function can be implemented by SAWR. Below, the optimisation problem and its solution are presented. Theorem 3.6.1 describes the result of the optimisation programme. We then show in section 3.6.3 that with the scoring rule from section 3.5.3 SAWR implements the optimal restricted direct mechanism. This concludes part (ii) of the analysis and verifies the guess restriction.

3.6.1   The Optimal Restricted Direct Mechanism

In a restricted direct mechanism that satisfies the guess restriction, the principal commits to a mechanism consisting of transfers $\boldsymbol{t}$, qualities $\boldsymbol{q}$, and an allocation decision rule $\boldsymbol{x}$ as functions of the agents' reports of types $\boldsymbol{\theta}$ and the winning agent's report of θ and ε in the second stage. Given the constraint from equation (3.66) and that the allocation decision is made before the realisation of the shock, the second stage transfers $\boldsymbol{t^2}$ and the qualities $\boldsymbol{q}$ may depend only on the winning agent's second stage report: $\widehat{\theta}, \widehat{\varepsilon}$. The allocation decision rule $\boldsymbol{x}$ and the first stage transfers $\boldsymbol{t^1}$ may only depend on the agents' type reports: $\widehat{\boldsymbol{\theta}}$.

The principal's problem is thus to design the variables $\boldsymbol{t^1} = \{t_1^1(\boldsymbol{\theta}), \cdots, t_N^1(\boldsymbol{\theta})\}$, $\boldsymbol{x}(\boldsymbol{\theta}) = \{x_1(\boldsymbol{\theta}), \cdots, x_N(\boldsymbol{\theta})\}$, $\boldsymbol{q}(\boldsymbol{\theta}, \boldsymbol{\varepsilon}) = \{q_1(\theta_1, \varepsilon), \cdots, q_N(\theta_N, \varepsilon)\}$, and $\boldsymbol{t^2}(\boldsymbol{\theta}, \boldsymbol{\varepsilon}) = \{t_1^2(\theta_1, \varepsilon), \cdots, t_N^2(\theta_N, \varepsilon)\}$ for all $\boldsymbol{\theta}, \varepsilon$ to maximise her utility subject to the restrictions that the agents wish to participate in the mechanism, truthfully reveal θ in the first stage, and truthfully reveal $\{\theta, \varepsilon\}$ in the second stage. The principal thus solves:

$$\max_{\boldsymbol{t^1}, \boldsymbol{t^2}, \boldsymbol{x}, \boldsymbol{q}} \sum_{i=1}^{N} E_{\boldsymbol{\theta}, \varepsilon}\left[\left(V(q_i(\theta_i) - t_i^2(\theta_i)\right)x_i(\boldsymbol{\theta}) - t_i^1(\boldsymbol{\theta})\right] \tag{3.67}$$

subject to first stage participation:

$$E_{\boldsymbol{\theta_{\text{-}i}},\varepsilon}\left[t_i^1(\theta_i,\boldsymbol{\theta_{\text{-}i}}) + \left(t_i^2(\theta_i,\varepsilon) - c(q_i(\theta_i,\varepsilon),\theta_i,\varepsilon)\right)x(\theta_i,\boldsymbol{\theta_{\text{-}i}})\right] \geq 0 \tag{3.68}$$

$(\forall i,\theta_i)$. As well as subject to second stage truthtelling from the guess restriction in equation (3.66):

$$\theta_i,\varepsilon = \underset{\widehat{\theta}_i,\widehat{\varepsilon}}{\operatorname{argmax}}\left\{t_i^2(\widehat{\theta}_i,\widehat{\varepsilon}) - c(q_i(\widehat{\theta}_i,\widehat{\varepsilon}),\theta_i,\varepsilon)\right\} \tag{3.69}$$

$(\forall i,\theta_i,\varepsilon)$. Note that given this strong second stage truthtelling constraint, the value $v_i(\widetilde{\theta}_i,\theta_i,\varepsilon)$ of an agent i with a first stage report of $\widetilde{\theta}_i$, a type θ_i and a shock ε of reaching the second stage:

$$v_i(\widetilde{\theta}_i,\theta_i,\varepsilon) = \underset{\widehat{\theta}_i,\widehat{\varepsilon}}{\max}\left\{t_i^2(\widehat{\theta}_i,\widehat{\varepsilon}) - c(q_i(\widehat{\theta}_i,\widehat{\varepsilon}),\theta_i,\varepsilon)\right\} \tag{3.70}$$

reduces to:

$$v_i(\widetilde{\theta}_i,\theta_i,\varepsilon) = t_i^2(\theta_i,\varepsilon) - c(q_i(\theta_i,\varepsilon),\theta_i,\varepsilon) \tag{3.71}$$

$(\forall i,\theta_i,\widetilde{\theta}_i)$. As such, unlike in section 3.4 and equation (3.7), the value of reaching the second stage is independent of the first stage report $\widetilde{\theta}_i$. The final constraint of first stage truthtelling thus becomes:

$$\theta_i \in \underset{\widetilde{\theta}_i}{\operatorname{argmax}}\, E_{\boldsymbol{\theta_{\text{-}i}},\varepsilon}\left[\left(t_i^2(\theta_i,\varepsilon) - c(q_i(\theta_i,\varepsilon),\theta_i,\varepsilon)\right)x_i(\widetilde{\theta}_i,\boldsymbol{\theta_{\text{-}i}}) + t_i^1(\widetilde{\theta}_i,\boldsymbol{\theta_{\text{-}i}})\right] \tag{3.72}$$

$(\forall i,\theta_i)$.

Since all agents are identical ex-ante and since we are in an i.i.d. environment, we can restrict attention to symmetric mechanisms in which agents are treated identically to simplify notation and drop the i subscript. Because of the assumed risk-neutrality, a change in second stage transfers can be compensated by an appropriate change in first stage transfers without affecting the agents' first stage behaviour. We may thus fix any one second stage transfer at an arbitrary constant and let the rest of the second stage transfers be determined by the second stage incentive constraint. The transfer we will want to fix depends on whether the quality policies are monotonic in the cost of producing quality. We cover both cases here.

Monotonicity

Assumption 7. Monotonicity 2. The optimal quality policies q^r, which form part of the solution to the restricted mechanism design problem, are non-increasing in the cost of producing quality. I.e. $q^r(\underline{\theta}, n) \leq q^r(\underline{\theta}, y) \leq q^r(\overline{\theta}, n) \leq q^r(\overline{\theta}, y)$.

This assumption is stronger than assumption 5. Any time it is satisfied, assumption 5 will also be. This assumption also excludes many reasonable model environments. It is used here since it allows us to show an analytical solution to the optimisation problem. Given monotonicity, we impose binding second stage participation for the unproductive type with the shock, i.e. $\{\theta = \overline{\theta}, \varepsilon = y\}$:

$$t^2(\overline{\theta}, y) = c(q(\overline{\theta}, y), \overline{\theta}, y). \tag{3.73}$$

By the convexity of the cost function, since the quality policies are monotonic in the cost of producing quality, second stage truthtelling is ensured in the least costly way for the principal when the following constraints are binding:

$$t^2(\overline{\theta}, n) = t^2(\overline{\theta}, y) + c(q(\overline{\theta}, n), \overline{\theta}, n) - c(q(\overline{\theta}, y), \overline{\theta}, n), \tag{3.74}$$

$$t^2(\underline{\theta}, y) = t^2(\overline{\theta}, n) + c(q(\underline{\theta}, y), \underline{\theta}, y) - c(q(\overline{\theta}, n), \underline{\theta}, y), \tag{3.75}$$

$$\text{and } t^2(\underline{\theta}, n) = t^2(\underline{\theta}, y) + c(q(\underline{\theta}, n), \underline{\theta}, n) - c(q(\underline{\theta}, y), \underline{\theta}, y). \tag{3.76}$$

The second stage transfers are thus completely determined by the agents' cost of producing different qualities q by equations (3.73) to (3.76). Let $W(\theta)$ denote the value of reaching the second stage for an agent of type θ. Given the previous analysis, we know that:

$$W(\theta = \overline{\theta}) = (1 - \delta)\Big(c(q(\overline{\theta}, y), \overline{\theta}, y) - c(q(\overline{\theta}, y), \overline{\theta}, n)\Big). \tag{3.77}$$

The value of the second stage is strictly positive in expectation over the shock even for the less productive type. Given binding first stage participation, agents of this type are willing to pay to the buyer in the first stage up to equal the amount of informational rent they can realise in the second stage (times the probability of getting there).

To get a formulation for $W(\underline{\theta})$ we need to use the first stage incentive constraints to get towards $E_{\boldsymbol{\theta}_{-i}}\left[t^1(n, \boldsymbol{\theta}_{-i})\right]$. To save on notation, we again let $t^1(\theta)$ denote $E_{\boldsymbol{\theta}_{-i}}\left[t^1(\theta, \boldsymbol{\theta}_{-i})\right]$ and $x(\theta)$ denote $E_{\boldsymbol{\theta}_{-i}}\left[x(\theta, \boldsymbol{\theta}_{-i})\right]$. The expected first stage transfers for an unproductive type agent $t^1(\overline{\theta})$ can then be found using the second stage transfers and the first stage participation constraint.

$$t^1(\overline{\theta}) = -E_\varepsilon\left[t^2(\overline{\theta}, \varepsilon) - c(q(\overline{\theta}, \varepsilon), \overline{\theta}, \varepsilon)\right]x(\overline{\theta}) \tag{3.78}$$

$$= -(1 - \delta)\left(c(q(\overline{\theta}, y), \overline{\theta}, y) - c(q(\overline{\theta}, y), \overline{\theta}, n)\right)x(\overline{\theta}) \tag{3.79}$$

$$= -W(\overline{\theta})x(\overline{\theta}). \tag{3.80}$$

Using equations (3.73) to (3.76) we can get an expression for $W(1)$ in terms of the cost functions:

$$W(\underline{\theta}) = (1 - \delta)\left(c(q(\underline{\theta}, y), \underline{\theta}, y) - c(q(\underline{\theta}, y), \underline{\theta}, n)\right) + c(q(\overline{\theta}, n), \overline{\theta}, n) - c(q(\overline{\theta}, n), \underline{\theta}, y)$$
$$+ c(q(\overline{\theta}, y), \overline{\theta}, y) - c(q(\overline{\theta}, y), \overline{\theta}, n). \tag{3.81}$$

The expected value of an agent of type n to reach the second stage is thus the cumulative informational rent he can extract from being able to be the other type and shock.

Interestingly, the first stage of the restricted optimal mechanism design problem resembles the standard setup of a basic optimal auction problem with private and independent values. We have N ex-ante identical agents with private information over their own type θ, which determines their valuation of the good $W(\theta)$ (i.e. the second stage). The principal's problem is to design transfer and allocation rules t^1 and x in order to maximise their expected revenue. However, crucially different from the simple private value auction, here for the agents the value of reaching the second stage W is endogenously determined by the quality policies q of the second stage. For the design of the optimal q, we thus have to evaluate their impact on the first and second stage simultaneously.

By the convexity of the cost function, the cheapest way for the principal to satisfy the first stage incentive constraints from equation (3.72) is to set it at binding for any type $\theta = \underline{\theta}$

agent pretending to be type y. The constraints from (3.72) thus reduce to:

$$W(\underline{\theta})x(\underline{\theta}) + t^1(\underline{\theta}) = W(\underline{\theta})x(\overline{\theta}) + t^1(\overline{\theta}). \tag{3.82}$$

Substituting for $t^1(\overline{\theta})$ from equation (3.80) and rearranging we get:

$$t^1(\underline{\theta}) = \left(W(\underline{\theta}) - W(\overline{\theta})\right)x(\overline{\theta}) - W(\underline{\theta})x(\underline{\theta}). \tag{3.83}$$

We are now ready to substitute the constraints into the optimisation programme and begin solving for the optimal x, t, and q. Rewriting the maximand from equation (3.67) we get:

$$\max_{t^1, t^2, x, q} N \cdot E_{\boldsymbol{\theta}, \varepsilon}\left[\left(V(q(\theta, \varepsilon)) - t^2(\theta, \varepsilon)\right) \cdot x(\boldsymbol{\theta}) - t^1(\boldsymbol{\theta})\right]. \tag{3.84}$$

Writing out the expectation over the representative agent's type θ and substituting for first stage transfers from equations (3.80) and (3.83) we get:

$$\max_{t^2, x, q} N \cdot E_{\theta, \boldsymbol{\theta_{-i}}, \varepsilon}\left[\left(V(q(\theta, \varepsilon)) - t^2(\theta, \varepsilon)\right)x(\theta) + W(\theta)x(\theta) - (1 - \gamma)\left(W(\underline{\theta}) - W(\overline{\theta})\right)x(\overline{\theta})\right]. \tag{3.85}$$

Using equations (3.73) to (3.76) to solve for $E_\varepsilon[t^2(\theta, \varepsilon)]$ we find:

$$E_\varepsilon\left[t^2(\theta, \varepsilon)\right] = W(\theta) + E_\varepsilon\left[c(q(\theta, \varepsilon), \theta, \varepsilon)\right] \tag{3.86}$$

$(\forall \theta)$. Substituting into the above (3.85) simplifies to:

$$\max_{x, q} N \cdot E_{\theta, \boldsymbol{\theta_{-i}}, \varepsilon}\left[\left(V(q(\theta, \varepsilon)) - c(q(\theta, \varepsilon), \theta, \varepsilon)\right)x(\theta) - (1 - \gamma)\left(W(\underline{\theta}) - W(\overline{\theta})\right)x(\overline{\theta})\right]. \tag{3.87}$$

From equation (3.77) and (3.81) we get:

$$W(\underline{\theta}) - W(\overline{\theta}) = (1 - \delta)\left(c(q(\underline{\theta}, y), \underline{\theta}, y) - c(q(\underline{\theta}, y), \underline{\theta}, n)\right) + \delta\left(c(q(\overline{\theta}, y), \overline{\theta}, y)\right.$$
$$\left. - c(q(\overline{\theta}, y), \overline{\theta}, n)\right) + \left(c(q(\overline{\theta}, n), \overline{\theta}, n) - c(q(\overline{\theta}, n), \underline{\theta}, y)\right). \tag{3.88}$$

Taking the first order conditions of the restricted optimal mechanism design programme

w.r.t. the optimal restricted qualities we get:

$$V'(q^r_{\underline{\theta}n}) = c_q(q^r_{\underline{\theta}n}, \underline{\theta}, n), \tag{3.89}$$

$$V'(q^r_{\underline{\theta}y}) = c_q(q^r_{\underline{\theta}y}, \underline{\theta}, y) + \frac{1-\delta}{\delta}\frac{x(\overline{\theta})}{x(\underline{\theta})}\Big(c_q(q^r_{\underline{\theta}y}, \underline{\theta}, y) - c_q(q^r_{\underline{\theta}y}, \underline{\theta}, n)\Big), \tag{3.90}$$

$$V'(q^r_{\overline{\theta}n}) = c_q(q^r_{\overline{\theta}n}, \overline{\theta}, n) + \frac{1-\gamma}{\gamma}\frac{1}{1-\delta}\Big(c_q(q^r_{\overline{\theta}n}, \overline{\theta}, n) - c_q(q^r_{\overline{\theta}n}, \underline{\theta}, y)\Big), \tag{3.91}$$

$$\text{and } V'(q^r_{\overline{\theta}y}) = c_q(q^r_{\overline{\theta}y}, \overline{\theta}, y) + \frac{1-\gamma}{\gamma}\Big(c_q(q^r_{\overline{\theta}y}, \overline{\theta}, y) - c_q(q^r_{\overline{\theta}y}, \overline{\theta}, n)\Big). \tag{3.92}$$

Here $q^r_{\theta\varepsilon}$ is the optimal quality of the restricted mechanism design programme given assumption 3.6.1.

These conditions for the optimal qualities of the restricted direct mechanism design problem are considerably different from the optimal quality conditions of the unrestricted programme. For the lowest cost type without the shock, the optimal qualities are undistorted in the restricted optimal mechanism. This is the same for the socially optimal qualities and the qualities of the optimal unrestricted mechanism. However, unlike the optimal mechanism, for the productive type with the cost shock the optimal qualities are now distorted if the fraction $\frac{x(\overline{\theta})}{x(\underline{\theta})}$ is positive, i.e. if there is a positive probability that an agent of the high type y is selected for the second stage. We have not yet said anything about the optimal allocation rules x, and in order to do so, we need first to notice that the distortion to the qualities $q^r_{\underline{\theta}y}$ is weakly positive. I.e., if there is a distortion, the qualities that solve equation (3.90) are smaller than the socially optimal qualities. Given this:

$$E_\varepsilon\Big[V(q^r_{\underline{\theta}\varepsilon}) - c(q(q^r_{\underline{\theta}\varepsilon}), \underline{\theta}, \varepsilon)\Big] > 0 \tag{3.93}$$

is immediate under the assumed functional forms. It is therefore optimal for the principal to set $x(\underline{\theta}, n) = 0.5$ and $x(\underline{\theta}, y) = 1 - x(\overline{\theta}, n)$. I.e., in the optimal restricted procurement mechanism it never pays for the principal to select no agent for the second stage if an agent of type n is available.

Unlike the low type, the high-cost type's optimal qualities are distorted with and without

the shock. Given the conditions for the optimal re qualities above,

$$E_\varepsilon\Big[V(q(q^r_{\underline{\theta}\varepsilon})) - c(q^r_{\underline{\theta}\varepsilon}, \underline{\theta}, \varepsilon)\Big] > E_\varepsilon\Big[V(q(q^r_{\overline{\theta}\varepsilon})) - c(q^r_{\overline{\theta}\varepsilon}, \overline{\theta}, \varepsilon)\Big] \tag{3.94}$$

always holds. Furthermore, $W(\underline{\theta}) - W(\overline{\theta})$ is strictly greater than zero. Therefore it is always better for the principal to select an agent of type n for the second stage and we get $x(\underline{\theta}, y) = 1$ and $x(\overline{\theta}, n) = 0$. Finally, from the optimisation programme we can derive the inequality:

$$\gamma E_\varepsilon\Big[V(q(q^r_{\overline{\theta}\varepsilon})) - c(q^r_{\overline{\theta}\varepsilon}, \overline{\theta}, \varepsilon)\Big] \geq (1 - \gamma)\Big(W(\underline{\theta}) - W(\overline{\theta})\Big). \tag{3.95}$$

If the inequality is satisfied, it is optimal for the principal to set $x(\overline{\theta}, y) = 0.5$. If it does not hold, she should set $x(\overline{\theta}, y) = 0$. Intuitively, the difference between $W(\underline{\theta})$ and $W(\overline{\theta})$ represents the informational rent the principal has to give up to agents of type n. If this rent exceeds the benefit of choosing an agent of type y when no agents of type n are available, the principal prefers in expectation to set $x(\overline{\theta}, y)$ to zero and thus limit the type rent of type n agents by making it worse for them to deviate their first stage report.

Proposition 3.6.1. *Given assumption 6, necessary and sufficient conditions for the optimal quality policies $q^r_{\theta\varepsilon}$ that form of the solution to the restricted mechanism design problem are given by the q that solve equations (3.89) to (3.92).*

Proof. The proof follows from the above analysis and the assumptions on the cost and utility functions. $\square$

3.6.2   A Reserve Score System for SAWR

The principal can use a reserve score system to improve her expected utility from the SAWR mechanism for any scoring rule. The SAWR implementation with a reserve score system is different from the one outlined in section 3.5.1 in the following points. In the timeline of events point 2 becomes:

2. The principal announces a second-scoreauction with scoring rule $S(p, q)$ and a reserve score system consisting of two reserve scores $\overline{S}^1$ and $\overline{S}^2$ with $\overline{S}^1 \leq \overline{S}^2$.

Furthermore, point 6 becomes:

6. If the winning agent's bid exceeds the first reserve score but the second-highest bid does not, i.e. if $S^1 > \overline{S}^1$ and $S^2 \leq \widetilde{S}^1$ are satisfied, the winning agent may choose a price p and quality q such that $S(p,q)$ matches the second reserve score, i.e. $S(p,q) \geq \overline{S}^2$. If this condition is not satisfied, the winning agent may choose p and q to match the score bid by the second-highest bidder, i.e. $S(p,q) \geq S_2$.

Finally points 1, 2-5, and 7 are the same as SAWR without a reserve score system in section 3.5.1. This peculiar reserve score system allows the principal to adjust SAWR such that even in a discrete type setting it sufficiently reduces the agents' rent to take into account that a deviation is made worse by the possibility of tie breakers. To understand how it works, we describe an equilibrium of SAWR with the reserve score system below.

Proposition 3.6.2. *A symmetric equilibrium of an efficient second-score SAWR mechanism[8] with discrete types and a reserve score system is one in which each winning firm of type θ, ε chooses:*

$$q(\theta, \varepsilon) = \underset{q}{\mathrm{argmax}} \, s(q) - c(q, \theta, \varepsilon) \tag{3.96}$$

and:

$$p(\theta, \varepsilon) = s(q(\theta, \varepsilon)) - \overline{S} \tag{3.97}$$

where $\overline{S}$ is the score required to be implemented by the agent in the second stage as described above and given the reserve scores below. There is an equilibrium of the first stage where each agent of type θ bids:

$$\widehat{S}(\theta) = E_\varepsilon \Big[s(q(\theta, \varepsilon)) - c(q(\theta, \varepsilon), \theta, \varepsilon) \Big] \tag{3.98}$$

and the principal chooses

$$\overline{S}^1 = \widehat{S}(\overline{\theta}) = E_S \Big[s(q(\overline{\theta}, \varepsilon)) - c(q(\overline{\theta}, \varepsilon), \overline{\theta}, \varepsilon) \Big] \tag{3.99}$$

---

[8]I.e. one where the contract is always allocated to an agent.

and

$$\overline{S}^2 = \widehat{S}(\overline{\theta}) + \frac{N-1}{N}\left(\widehat{S}(\underline{\theta}) - \widehat{S}(\overline{\theta})\right).$$ (3.100)

Proof. The second stage strategies are equilibrium strategies by the argument outlined above the proposition. To see the first stage strategies first note that neither type has any incentive to bid a different score. This is mostly a straightforward extension of Vickrey (1961) and Che (1993). For instance, a productive type has no incentive to bid higher as although he wins the auction 100% of the time, if there is another productive agent he does not make any profits and if there is no other productive agent he makes the same profits and would have won either way. He also has no incentive to bid lower as he will still be asked to match the same score $\overline{S}^2$ if no other agent is the productive type. However, note that bidding as described here is no longer a dominant strategy for the agents. Finally, to see that type $\theta = \underline{\theta}$ does not want to deviate and bid $\widehat{S}(\overline{\theta})$ note that this would lead to winning the contract $\frac{1}{N}\gamma^{N-1}$ of the time and when winning realising profits of $\widehat{S}(\underline{\theta}) - \widehat{S}(\overline{\theta})$. If the agent bids $\widehat{S}(\underline{\theta})$ he wins γ^{N-1} of the time and when winning realises profits $\left(1 - \frac{N-1}{N}\right)\left(\widehat{S}(\underline{\theta}) - \widehat{S}(\overline{\theta})\right)$. Thus the expected profits are exactly the same, and this is not a profitable deviation for the risk-neutral agent. $\qquad\square$

3.6.3 Equilibrium Implementation by SAWR

The final step to showing that we have identified the optimal scoring auction with renegotiation mechanism is to show that SAWR can implement the optimal restricted mechanism in equilibrium. To this end, consider a scoring auction as in section 3.6.2 where the scoring rule satisfies assumption 6 and equations (3.46) to (3.49) at the qualities given by equations (3.89) to (3.92). We know by the analysis above that the expected utility for the principal of this scoring auction when the inequality in equation (3.95) is satisfied will be as in equation

(3.50), except that $E_{\theta\varepsilon}[\pi^{Sc}(\theta)]$ will now be given by:

$$E_{\theta\varepsilon}[\pi^{Sc}(\theta)] = (1-\gamma)\gamma^{N-1}\left(1 - \frac{N-1}{N}\right)\left((1-\delta)\Big(c(q_{\underline{\theta}y}^r),\underline{\theta},y) - c(q_{\underline{\theta}y}^r,\underline{\theta},n)\Big)\right.$$
$$\left. + \delta\Big(c(q_{\overline{\theta}y}^r,\overline{\theta},y) - c(q_{\overline{\theta}y}^r,\overline{\theta},n)\Big) + c(q_{\overline{\theta}n}^r,\overline{\theta},n) - c(q_{\overline{\theta}n}^r,\underline{\theta},y)\right). \tag{3.101}$$

In the optimal restricted mechanism, the principal earns

$$N \cdot E_{\theta,\boldsymbol{\theta}_{-i},\varepsilon}\left[\Big(V(q_{\theta\varepsilon}^r) - c(q_{\theta\varepsilon}^r,\theta,\varepsilon)\Big)x(\theta) - (1-\gamma)\Big(W(\underline{\theta}) - W(\overline{\theta})\Big)x(\overline{\theta})\right]. \tag{3.102}$$

where by equation (3.88), $W(\underline{\theta}) - W(\overline{\theta})$ are:

$$W(\underline{\theta}) - W(\overline{\theta}) = (1-\delta)\Big(c(q_{\underline{\theta}y}^r,\underline{\theta},y) - c(q_{\underline{\theta}y}^r,\underline{\theta},n)\Big) + \delta\Big(c(q_{\overline{\theta}y}^r,\overline{\theta},y)$$
$$- c(q_{\overline{\theta}y}^r,\overline{\theta},n)\Big) + \Big(c(q_{\overline{\theta}n}^r,\overline{\theta},n) - c(q_{\overline{\theta}n}^r,\underline{\theta},y)\Big). \tag{3.103}$$

We can thus rewrite equation (3.101) as:

$$E_{\theta\varepsilon}[\pi^{Sc}(\theta)] = (1-\gamma)\gamma^{N-1}\left(1 - \frac{N-1}{N}\right)\Big(W(\underline{\theta}) - W(\overline{\theta})\Big). \tag{3.104}$$

The difference in expected utility between SAWR and the optimal restricted mechanism is thus zero if the condition

$$N(1-\gamma)\gamma^{N-1}\left(1 - \frac{N-1}{N}\right) = N(1-\gamma)x(\overline{\theta}) \tag{3.105}$$

is satisfied. For any allocative efficient restricted optimal mechanism (where condition (3.95) holds), the optimal allocation rule $x(\overline{\theta})$ is given by:

$$x(\overline{\theta}) = \text{Prob}(\boldsymbol{\theta}_{\boldsymbol{-i}} = \overline{\boldsymbol{\theta}}^{\boldsymbol{N-1}}) \cdot \frac{1}{N} \tag{3.106}$$

$$= \gamma^{N-1}\frac{1}{N}. \tag{3.107}$$

Equation (3.105) thus holds at equality. Since the two mechanisms lead to the same expected utility, and since they implement the same qualities and allocation rules, SAWR implements the optimal restricted mechanism.

Proposition 3.6.3. *A SAWR mechanism with a reserve score system as in proposition 3.6.2, the qualities of the restricted optimal mechanism in section 3.6, and a scoring rule that implements these qualities with the lowest possible total transfers as described by equations (3.73) to (3.76) is the best possible SAWR mechanism subject to assumption 3.6.1.*

Proof. The proof is given by the above analysis. Lemma 3.6.1 shows that the guess restriction does not exclude any possible SAWR mechanisms. Proposition 3.6.1 describes the optimal mechanism given the guess restrictio, and the above analysis shows that SAWR with a reserve score system as in proposition 3.6.2 implements the optimal restricted mechanism. □

3.6.4 Numerical Example

In the previous sections, we derived the optimal scoring auction with renegotiation. To determine how much is lost by going from the optimal mechanism to this optimal SAWR mechanism, we can compare the expected payoffs to the principal between the two mechanisms. Recall that the payoffs for the NHS from the different mechanisms are given in section 3.5.1. We summarise these as well as the payoffs of using the optimal SAWR mechanism in table 3.1.

As is the case with the naïve SAWR and the one with the qualities of the optimal mechanism previously, unsurprisingly the optimal SAWR leads to a loss in expected utility for the buyer relative to the optimal mechanism. We can also see that the optimal SAWR mechanism is only able to recoup one third (half) of the loss in EU of the SAWR mechanism with the OM qualities (naïve SAWR). The optimal SAWR mechanism only allocates 76% of the rent of the naïve SAWR and 88% of the rent of SAWR with OM qualities. This is equivalent to a 24% and 12% decrease in profitability for the agents. Figure 3.3 depicts graphically how the social surplus, the expected utility of the buyer, and the total rent extracted by the agents varies across the mechanisms. From left to right, the optimal mechanism (OM) has the highest expected buyer utility and the lowest rent paid to the agents. With a naïve

Mechanism	Expected Utility	Total Rent	Social Surplus	% loss
Optimal Mechanism	1.279	0.030	1.309	0 %
SAWR with OM q	1.226	0.083	1.309	4.15 %
- *with reserve score*	1.268	0.042	1.309	0.89 %
Naïve SAWR	1.217	0.098	1.315	4.88 %
- *with reserve score*	1.266	0.049	1.315	1.05 %
Optimal SAWR	1.234	0.074	1.308	3.49 %
- *with reserve score*	1.271	0.037	1.308	0.59 %

Table 3.1 Expected utilities

scoring auction and no reserve score system we have the highest social surplus, the lowest buyer utility, and consequently the highest total rent. The scoring auction that comes closest to implementing the optimal mechanism is the scoring auction with qualities from the restricted optimal mechanism and a reserve score system. This is depicted at the very right.

3.7 Conclusion

In this chapter, we derived the optimal mechanism for a dynamic procurement problem in which there is uncertainty about the future cost of the project. This is a kind of environment that we see commonly in the practice of procurement, as we saw in chapter 1. We showed that one of the most prominent mechanisms used to account for the multidimensionality of procurement problems is not optimal from the buyer's perspective if renegotiation is used to respond to the cost-uncertainty. Following this, we introduced a methodology that allowed us to describe the optimal scoring auction with renegotiation as the best of a sub-optimal
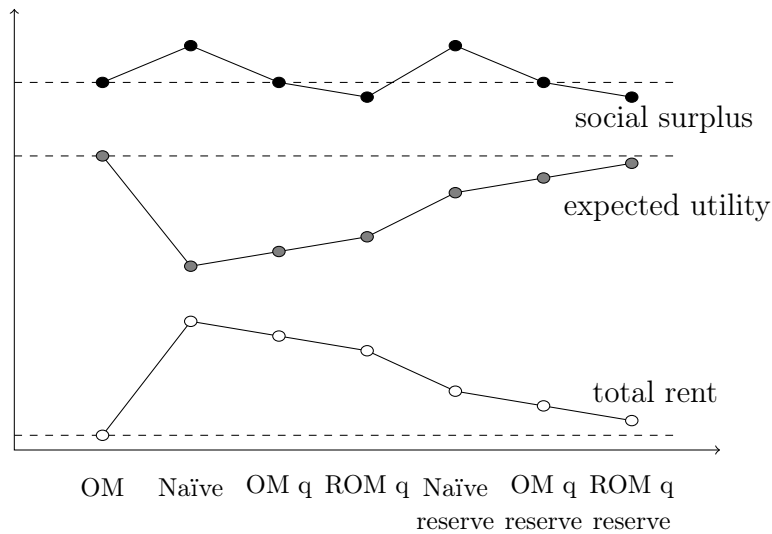
Figure 3.3 Social Surplus, Expected Utilities, and Agent Rent in the different mechanisms (not to scale)

class of mechanisms and the scoring rule associated with it.

In the numerical example of the NHS, the relative percentage gains in utility for the NHS from implementing the optimal SAWR over a naïve or sub-optimal one are small. Largely, this is because the rent extracted by the agents is low relative to the utility gain from the procurer even with a simple auction mechanism. Since the changes are considerably larger from the perspective of the agents' profitability, and since these numbers were arbitrarily chosen with no consideration to realism, this cannot be interpreted as strong evidence that the gains are too small for an attempt to implement the optimum to be not worth it. However, there are several important points we can take away.

First, the loss from not implementing a reserve price system was consistently larger than the loss from using the wrong scoring rule. This, however, will depend on the number of agents taking part in the auction and will not necessarily hold for auctions with more bidders. Second, the private gain for the buyer from the optimal mechanism or the optimal scoring auction, relative to the loss in social surplus is very small. This is particularly the case when the qualities are not calculated correctly, for example, in the scoring auction with the qualities from the optimal mechanism. In order to appropriate qualities and scoring rules,

we needed to impose strong assumptions on functional forms. Additionally, we are assuming that the buyer knows these forms as well as the distributions of types and shocks. If we consider these factors, the "naïve" scoring auction with renegotiation becomes a robust and well-performing option. It does not require the buyer to know anything other than her own preferences, it achieves the highest social surplus possible, and it considerably simplifies the auction as well as the renegotiation procedure.

Appendix

B.1 Proof of Lemma 3.6.1

The bundle p_i^*, q_i^* chosen by any winning agent i in equilibrium of the second stage of SAWR are the p, q that maximise the agent's utility given the scoring rule $S(p, q)$, subject to the restriction that $S(p_i^*, q_i^*) \geq \overline{S}$ where $\overline{S}$ is some score to be satisfied by the winning agent. To this end, the agent may choose *any* p, q as long as they achieve the desired score. In other words:

$$p_i^*, q_i^* \in \underset{p,q}{\mathrm{argmax}} \left\{ p - c(q, \theta, \varepsilon) \right\} \tag{B.1}$$

$(\forall i, \theta, \varepsilon)$ s.t. $S(p, q) \geq \overline{S}$. The agent's utility optimising strategy is to choose the pair of p and q that maximise the difference between the price and the cost of producing the quality. With more general notation, the second stage strategy ϕ_i^2 of agent i thus depends on the agent's cost type θ_i, the shock ε, as well as his first stage strategy ϕ_i^1, which includes the score he bids in the first stage. We can thus equivalently write equation (B.1) as:

$$\phi_i^{2*}(\theta_i, \varepsilon | \phi_i^1(\theta_i)) \in \underset{\widetilde{\phi}_i^2}{\mathrm{argmax}} \left\{ \pi_i \left(g \left(\phi_i^1(\theta_i), \widetilde{\phi}_i^2, \boldsymbol{\phi_{\text{-}i}^{1*}}(\boldsymbol{\theta_{\text{-}i}}) \right), \theta_i, \varepsilon \right) \right\} \tag{B.2}$$

$(\forall i, \theta, \varepsilon, \phi_i^1(\theta_i), \boldsymbol{\phi_{\text{-}i}^{1*}}(\boldsymbol{\theta_{\text{-}i}}))$. The agent's equilibrium second stage strategy in SAWR, denoted $\phi_i^{2*}(\theta_i, \varepsilon | \phi_i^1(\theta_i))$, is thus the strategy that maximises his profits π_i as a function of the outcome g and his type θ_i, ε. The outcome function g here just specifies the price and quality for the agent given his second stage strategy, first stage strategy, as well as the other agent's

equilibrium strategy. Note that in SAWR, by winning the auction and being asked to implement a certain score, the agent effectively learns the other agent's type by the second stage (given he is following the equilibrium strategy).

In the first stage of SAWR, an agent bids p and q such as to achieve the score S^* that maximises his expected profits. Through the probability of winning and the second stage renegotiation mechanism, this choice is influenced by his type θ_i, his expectation over the shock ε, and his expectation over the other agents' types $\boldsymbol{\theta_{-i}}$ and equilibrium bidding behaviour $\boldsymbol{\phi^*_{-i}}(\boldsymbol{\theta_{-i}})$. We can summarise this strategy in the above notation as:

$$\phi_i^{1*}(\theta_i) \in \underset{\widetilde{\phi}^1}{\operatorname{argmax}} \, E_{\varepsilon, \boldsymbol{\theta_{-i}}} \left[\pi_i \left(g\left(\widetilde{\phi}^1, \phi_i^{2*}(\theta_i, \varepsilon | \widetilde{\phi}^1), \boldsymbol{\phi^{1*}_{-i}}(\boldsymbol{\theta_{-i}}) \right), \theta_i, \varepsilon \right) \right] \tag{B.3}$$

$(\forall i, \theta_i)$. Here ϕ_i^{2*} is as defined above. So agents in the first stage choose the score they bid optimally to maximise their expected utility, knowing that in the second stage they will behave optimally whatever their first stage bid.

Let $f(\cdot)$ be the equilibrium outcome function implemented by a SAWR mechanism in perfect Bayesian Nash equilibrium. As shown in chapter 2, by the revelation principle, this function can be implemented by a direct mechanism Γ^{DM} of the form used in section 3.4 where agents directly reveal their type θ first and the winning agent reveals the shock ε. However, to show lemma 3.6.1 holds, we need to show that this social choice function can also be implemented with a direct mechanism subject to the added truthtelling constraint from equation (3.66).

To this end, assume an agent i of private information type θ_i and shock ε reaches the second stage with an arbitrary first stage strategy $\widetilde{\phi}_i^1$. The principal may ask the agent to report a cost parameter and a shock bundle $\widehat{\theta_i, \varepsilon}$, figure out $\phi_i^{2*}(\widehat{\theta_i, \varepsilon} | \widetilde{\phi}_i^1)$ and assign the agent $g(\widetilde{\phi}_i^1, \phi_i^{2*}(\widehat{\theta_i, \varepsilon} | \widetilde{\phi}_i^1), \boldsymbol{\phi_{-i}^*})$. Since the deviation $\phi_i^{2*}(\widetilde{\theta_i, \varepsilon} | \widetilde{\phi}^1), \widetilde{\theta_i, \varepsilon} \neq \theta_i, \varepsilon$ was available as a deviation $\widetilde{\phi}_i^{2*}(\theta_i, \varepsilon | \widetilde{\phi}^1)$ in equation (B.2)[1], the agent will reveal θ_i, ε truthfully in equilib-

---

[1]In any second stage of SAWR, an agent may behave as would be optimal for an agent of any different θ_i, ε given the first stage score $\widetilde{\phi}^1$.

rium of this new second stage. As in appendix A.1 of chapter 2, we need to account for a possible difference in the expectation over $\boldsymbol{\theta_{-i}}$ between the direct mechanism and the SAWR mechanism since the winning agent may learn about the other agents' types by winning. To do this, we can assume in the same way as in chapter 2 that in the direct mechanism the principal reveals to the winning agent whatever she would have learnt in SAWR, i.e. she reveals to the winning agent the second-highest agent's type. In the more general language introduced above:

$$\theta_i, \varepsilon \in \underset{\widetilde{\theta_i, \varepsilon}}{\operatorname{argmax}} \left\{ \pi_i \left(g\left(\widetilde{\phi}_i^1, \phi_i^{2*}(\widetilde{\theta_i, \varepsilon}|\widetilde{\phi}_i^1), \boldsymbol{\phi_{-i}^*}(\boldsymbol{\theta_{-i}}) \right), \theta_i, \varepsilon \right) \right\} \tag{B.4}$$

$(\forall i, \theta_i, \varepsilon, \widetilde{\phi}^1)$. From equation (B.3), $\phi_i^{1*}(\theta_i)$ will still be a first stage equilibrium strategy given the new second stage if the first stage remains unchanged. The principal may then replace the first stage with an alternative in which the agents are asked to reveal a type $\widehat{\theta}_i$ and are allocated a strategy $\phi_i^{1*}(\widehat{\theta}_i)$ according to their type reports. Since any deviation $\phi_i^{1*}(\widetilde{\theta}_i)$ was available as part of $\widetilde{\phi}_i^1(\theta_i)$ in equation (B.3), truthtelling is an equilibrium strategy for any agent in the new mechanism. In other words:

$$\theta_i \in \underset{\widetilde{T}_i}{\operatorname{argmax}} \, E_{\boldsymbol{\theta_{-i}}, \varepsilon} \left[\pi_i \left(g\left(\phi_i^1(\widetilde{\theta}_i), \phi_i^{2*}(\theta_i, \varepsilon|\phi_i^1(\widetilde{\theta}_i)), \boldsymbol{\phi_{-i}^*}(\boldsymbol{\theta_{-i}}) \right), \theta_i, \varepsilon \right) \right] \tag{B.5}$$

$(\forall i, \theta_i)$. This holds as long as the buyer commits to not using the revealed θ_i to infer the agent's type for the second stage.[2] Conditions (B.4) and (B.5) are precisely the conditions for the social choice function $f(\cdot)$ of the SAWR mechanism to be truthfully implementable in Bayesian Nash equilibrium of a direct mechanism subject to the additional constraint given by equation (3.66), which concludes the proof.

---

[2] Of course this is not optimal for the buyer. This is not surprising since SAWR is not an optimal mechanism, as it does not make optimal use of the revealed private information of the agents.

REFERENCES

Armstrong, Mark and David EM Sappington (2007). "Recent developments in the theory of regulation". In: *Handbook of industrial organization* 3, pp. 1557–1700.

Arve, Malin and David Martimort (2016). "Dynamic Procurement under Uncertainty: Optimal Design and Implications for Incomplete Contracts". In: *The American Economic Review* 106.11, pp. 3238–3274.

Asker, John and Estelle Cantillon (2005). "Optimal procurement when both price and quality matter". In:

— (2008). "Properties of scoring auctions". In: *The RAND Journal of Economics* 39.1, pp. 69–85.

— (2010). "Procurement when price and quality matter". In: *The Rand journal of economics* 41.1, pp. 1–34.

Bajari, Patrick, Stephanie Houghton, and Steven Tadelis (2014). "Bidding for incomplete contracts: An empirical analysis of adaptation costs". In: *American Economic Review* 104.4, pp. 1288–1319.

Bajari, Patrick, Robert McMillan, and Steven Tadelis (2008). "Auctions versus negotiations in procurement: an empirical analysis". In: *The Journal of Law, Economics, & Organization* 25.2, pp. 372–399.

Bajari, Patrick, Steven Tadelis, et al. (2001). "Incentives versus transaction costs: A theory of procurement contracts". In: *Rand journal of Economics* 32.3, pp. 387–407.

Bergemann, Dirk and Maher Said (2010). "Dynamic auctions". In: *Wiley Encyclopedia of Operations Research and Management Science*.

Branco, Fernando (1997). "The design of multidimensional auctions". In: *The RAND Journal of Economics*, pp. 63–81.

Bucciol, Alessandro, Riccardo Camboni, and Paola Valbonesi (2020). *Purchasing medical devices: the role of buyers' ability and discretion*.

Bulow, Jeremy and Paul Klemperer (1994). *Auctions vs. negotiations*. Tech. rep. National Bureau of Economic Research.

Burguet, Roberto and Yeon-Koo Che (2004). "Competitive procurement with corruption". In: *RAND Journal of Economics*, pp. 50–68.

Butler, Jeffrey V et al. (2020). "Past performance and entry in procurement: An experimental investigation". In: *Journal of Economic Behavior & Organization* 173, pp. 179–195.

Calzolari, Giacomo and Giancarlo Spagnolo (2009). "Relational contracts and competitive screening". In:

Camboni, Riccardo, Paola Valbonesi, et al. (2019). *Incumbent and entrant bidding in scoring rule auctions: A study on Italian canteen services*. Tech. rep. Dipartimento di Scienze Economiche" Marco Fanno".

Celentani, Marco and Juan-José Ganuza (2002). "Corruption and competition in procurement". In: *European Economic Review* 46.7, pp. 1273–1303.

Chang, Wei-Shiun, Timothy C Salmon, and Krista J Saral (2016). "Procurement auctions with renegotiation and wealth constraints". In: *Economic Inquiry* 54.3, pp. 1684–1704.

Chao, Hung-Po and Robert Wilson (2002). "Multi-dimensional procurement auctions for power reserves: Robust incentive-compatible scoring and settlement rules". In: *Journal of Regulatory Economics* 22.2, pp. 161–183.

Chassang, Sylvain and Juan Ortner (2019). "Collusion in auctions with constrained bids: Theory and evidence from public procurement". In: *Journal of Political Economy* 127.5, pp. 000–000.

Che, Yeon-Koo (1993). "Design competition through multidimensional auctions". In: *The RAND Journal of Economics*, pp. 668–680.

Dalen, Dag Morten, Espen R Moen, and Christian Riis (2006). "Contract renewal and incentives in public procurement". In: *International Journal of Industrial Organization* 24.2, pp. 269–285.

Decarolis, Francesco et al. (2018). *Bureaucratic competence and procurement outcomes*. Tech. rep. National Bureau of Economic Research.

Dellarocas, Chrysanthos, Federico Dini, and Giancarlo Spagnolo (2006). "18 Designing reputation mechanisms". In: *Handbook of Procurement*, p. 446.

Dewatripont, Mathias (1989). "Renegotiation and information revelation over time: The case of optimal labor contracts". In: *The Quarterly Journal of Economics* 104.3, pp. 589–619.

Dimitri, Nicola, Gustavo Piga, and Giancarlo Spagnolo (2006). "Introduction". In: *Handbook of Procurement*. Cambridge University Press, pp. 3–13. DOI: 10.1017/CBO9780511492556.002.

Doni, Nicola (2006). "The importance of reputation in awarding public contracts". In: *Annals of Public and Cooperative Economics* 77.4, pp. 401–429.

Eső, Péter and Balazs Szentes (2007). "Optimal information disclosure in auctions and the handicap auction". In: *The Review of Economic Studies* 74.3, pp. 705–731.

Estache, Antonio et al. (2009). "Multidimensionality and renegotiation: Evidence from transport sector public-private-partnership transactions in Latin America". In: *Review of Industrial Organization* 35.1-2, p. 41.

Farrell, Joseph and Paul Klemperer (2007). "Coordination and lock-in: Competition with switching costs and network effects". In: *Handbook of industrial organization* 3, pp. 1967–2072.

Ganuza, Juan-José (2007). "Competition and cost overruns in procurement". In: *The Journal of Industrial Economics* 55.4, pp. 633–660.

Guasch, J Luis, Jean-Jacques Laffont, and Stéphane Straub (2006). "Renegotiation of concession contracts: a theoretical approach". In: *Review of Industrial organization* 29.1-2, pp. 55–73.

Guasch, J Luis, Stéphane Straub, and Jean-Jacques Laffont (2003). *Renegotiation of concession contracts in Latin America*. The World Bank.

Guerrieri, Veronica and Robert Shimer (2014). "Dynamic adverse selection: A theory of illiquidity, fire sales, and flight to quality". In: *American Economic Review* 104.7, pp. 1875–1908.

Hart, Oliver, Andrei Shleifer, and Robert W Vishny (1997). "The proper scope of government: theory and an application to prisons". In: *The Quarterly Journal of Economics* 112.4, pp. 1127–1161.

Hendel, Igal, Alessandro Lizzeri, and Marciano Siniscalchi (2005). "Efficient sorting in a dynamic adverse-selection model". In: *The Review of Economic Studies* 72.2, pp. 467–497.

Herweg, Fabian and Klaus M Schmidt (2017). "Auctions versus negotiations: the effects of inefficient renegotiation". In: *The RAND Journal of Economics* 48.3, pp. 647–672.

— (2018). "Procurement with Unforeseen Contingencies". In: *Working Paper*. URL: https://editorialexpress.com/cgi-bin/conference/download.cgi?db_name=EARIE45&paper_id=562.

Herweg, Fabian and Marco A Schwarz (2018). "Optimal Cost Overruns: Procurement Auctions with Renegotiation". In: *International Economic Review* 59.4, pp. 1995–2021.

Ishii, Rieko (2008). "Collusion in repeated procurement auction: a study of a paving market in Japan". In:

Jofre-Bonet, Mireia and Martin Pesendorfer (2000). "Bidding behavior in a repeated procurement auction: A summary". In: *European Economic Review* 44.4-6, pp. 1006–1020.

Johnson, Terence R (2013). *Procurement with adverse selection and dynamic moral hazard.* Tech. rep. working paper.

Jung, Hojin et al. (2019). "Strategic bidding and contract renegotiation". In: *International Economic Review* 60.2, pp. 801–820.

Klemperer, Paul (1999). "Auction theory: A guide to the literature". In: *Journal of economic surveys* 13.3, pp. 227–286.

Krishna, Vijay (2009). *Auction theory.* Academic press.

Laffont, Jean-Jacques and Jean Tirole (1987). "Auctioning incentive contracts". In: *Journal of Political Economy* 95.5, pp. 921–937.

— (1990). "Adverse selection and renegotiation in procurement". In: *The Review of Economic Studies* 57.4, pp. 597–625.

— (1991). "Auction design and favoritism". In: *International Journal of Industrial Organization* 9.1, pp. 9–42.

Lewis, Tracy R and Huseyin Yildirim (2005). "Managing switching costs in multiperiod procurements with strategic buyers". In: *International Economic Review* 46.4, pp. 1233–1269.

— (2006). "Managing Dynamic Procurement". In: *N., Dimitri, G., Piga, G., Spagnolo (Eds.), Handbook of Procurement*, pp. 433–445.

Manelli, Alejandro M and Daniel R Vincent (1995). "Optimal procurement mechanisms". In: *Econometrica: Journal of the Econometric Society*, pp. 591–620.

Mas-Colell, Andreu, Michael Dennis Whinston, Jerry R Green, et al. (1995). *Microeconomic theory.* Vol. 1. Oxford university press New York.

McAfee, R Preston and John McMillan (1987). "Auctions with entry". In: *Economics Letters* 23.4, pp. 343–347.

Myerson, Roger B (1981). "Optimal auction design". In: *Mathematics of operations research* 6.1, pp. 58–73.

OECD (2015). *Government at a Glance.* OECD publishing. URL: stats.oecd.org.

Pai, Mallesh M and Rakesh Vohra (2013). "Optimal dynamic auctions and simple index rules". In: *Mathematics of Operations Research* 38.4, pp. 682–697.

Pavan, Alessandro, Ilya Segal, and Juuso Toikka (2014). "Dynamic mechanism design: A myersonian approach". In: *Econometrica* 82.2, pp. 601–653.

Quint, Daniel and Kenneth Hendricks (2018). "A theory of indicative bidding". In: *American Economic Journal: Microeconomics* 10.2, pp. 118–51.

Riley, John G and William F Samuelson (1981). "Optimal auctions". In: *The American Economic Review* 71.3, pp. 381–392.

Riordan, Michael H and David EM Sappington (1987a). "Awarding monopoly franchises". In: *The American Economic Review*, pp. 375–387.

— (1987b). "Information, incentives, and organizational mode". In: *The Quarterly Journal of Economics* 102.2, pp. 243–263.

Rose-Ackerman, Susan (1996). "Altruism, nonprofits, and economic theory". In: *Journal of economic literature* 34.2, pp. 701–728.

Thomas, Charles J (2010). "Information revelation and buyer profits in repeated procurement competition". In: *The Journal of Industrial Economics* 58.1, pp. 79–105.

Tirole, Jean (1986). "Procurement and renegotiation". In: *Journal of Political Economy* 94.2, pp. 235–259.

Vickrey, William (1961). "Counterspeculation, auctions, and competitive sealed tenders". In: *The Journal of finance* 16.1, pp. 8–37.

Wang, Ruqu (2000). "Bidding and renegotiation in procurement auctions". In: *European Economic Review* 44.8, pp. 1577–1597.