

Essays in Market and Mechanism Design



Inauguraldissertation zur Erlangung des akademischen Grades
eines Doktors der Wirtschaftswissenschaften der Universität
Mannheim

Marco Reuter

Frühjahrs-/Sommersemester 2023

Abteilungssprecher: Prof. Klaus Adam, Ph.D.
Referent: Prof. Dr. Thomas Troger
Korreferent: Prof. Dr. Vitali Gretschko
Datum: 08. Mai, 2023

This dissertation is the result of my own work and no other sources or means, except the ones listed, have been employed.

*Dành cho vĩ côa m¹nh
- em Tróc xinh -*

Contents

Preface	vii
Acknowledgements	ix
1 The Value of Decentralization Using the Blockchain	1
1.1 Introduction	1
1.2 Model	7
1.2.1 Centralized Governance	9
1.2.2 Decentralized Governance	15
1.2.3 Optimal Governance	17
1.2.4 Welfare	19
1.3 Discussion	20
1.3.1 Airdrops: Decentralizing at a Later Time	20
1.3.2 Equilibrium Multiplicity	21
1.4 Conclusion	22
1.5 Appendix	23
1.5.1 Explanation of Blockchain, Smart Contracts and the Creation of Commitment	23
1.5.2 Myopic Revenue Maximization	26
1.5.3 Sufficient conditions for $\bar{\mu}$ to be well-defined	29
1.5.4 Proof of Proposition 1	33
1.5.5 Proof of Lemma 2	36
1.5.6 Proof of Proposition 2	38
1.5.7 Proof of Lemma 3	38
1.5.8 Proof of Proposition 3	38
1.5.9 Proof of equilibrium of section 1.3.2 and proof of proposition 5	40
1.5.10 Proof of Lemma 4	41
1.5.11 Proof of proposition 4	41
1.5.12 Extension: Pre-commitment to revenue sharing path in decentral- ized governance	43

2	Mechanism Design for Unequal Societies	44
2.1	Introduction	44
2.2	Related Literature	49
2.3	Framework	51
2.4	Optimal mechanisms	55
2.4.1	No-subsidy constraints	55
2.4.2	Ex ante budget balance	58
2.5	Implementation via auctions	62
2.5.1	No subsidy constraints	62
2.5.2	Ex ante budget balance	64
2.6	Numerical illustrations	66
2.7	Conclusion	68
2.8	Appendix for Mechanism Design for Unequal Societies	69
2.8.1	Proof of Proposition 1	69
2.8.2	Proof of Lemma 1	69
2.8.3	Proof of Proposition 2	70
2.8.4	Proof of Proposition 3	73
2.8.5	Proof of Remark 1	74
2.8.6	Proof of Corollary 3	77
2.8.7	Proof of Proposition 4	78
2.8.8	Proof of Proposition 5	78
2.8.9	Proof of Proposition 6	80
2.8.10	Proof of Proposition 7	85
2.8.11	Calculations of Bidding Subsidies for the Specific Examples	88
2.8.12	Derivation of the Uniform v^K , Pareto v^M Example	90
3	Revenue Maximization with Partially Verifiable Information	93
3.1	Introduction	93
3.2	Model	98
3.2.1	Description	98
3.2.2	Mechanisms and the Revelation Principle for Partially Verifiable Types	101
3.2.3	Incentive Compatibility and Expected Revenue	103
3.2.4	Deriving the Optimal Mechanism	105
3.3	Two-Stage Implementation & A Right to Remain Anonymous	108
3.3.1	A Two-Stage Implementation	108
3.3.2	A Right to Remain Anonymous	111
3.4	Conclusion	113
3.5	Appendix for Revenue Maximization with Partially Verifiable Information	114

3.5.1	Explanation of the Initial and Extended Environment	114
3.5.2	Proof of Proposition 1	115
3.5.3	Proof of Proposition 2	116
3.5.4	Proof of Lemma 1	118
3.5.5	Proof of Proposition 3	120
3.5.6	Proof of Lemma 2	121
3.5.7	Proof of Lemma 3	121
3.5.8	Proof of Lemma 4	122
3.5.9	Proof of Lemma 5	123
3.5.10	Proof of Proposition 4	123
3.5.11	Proof of Proposition 5	124
3.5.12	Proof that Non-Decreasing Virtual Valuations are Preserved Under Mixing	125

Preface

Every day, we interact with each other in various markets and economic mechanisms. The rules that shape these markets and mechanisms shape our economic interactions. In particular, rules create incentives that govern economic interactions. The design of rules then allows us to shape economic incentives to pursue different market outcomes. This thesis is dedicated to studying the design of markets and mechanisms, their rules, incentives and goals in three applications.

Blockchains: Blockchains are a new and rapidly developing technology. Most prominently, they are used in cryptocurrencies. They have been experiencing alternating periods of hype and disappointment. Proponents of the technology argue that it has the potential to change economic interactions fundamentally, while critics see little economic value. In short, there does not seem to be a consensus whether, and in which situations, Blockchains bring any value to economic interactions. Recently, an exciting new strand of literature has developed that examines this question in detail. I contribute to this literature by highlighting the value that commitment created by a blockchain can have for a particularly important application in this day and age: creating a platform with network effects that locks in its users.

In the first chapter, I develop a model of an entrepreneur, who can create a network for her users. She can decide to retain control of the network with centralized implementation through a regular company, or surrender control over the network with a decentralized implementation through the blockchain. Users that join the network are subject to a locked-in effect. As the main finding, I show that a decentralized implementation of the network is (i) preferred by the entrepreneur and (ii) a Pareto improvement, if and only if the size of the locked-in effect is sufficiently large.

Inequality: In the second chapter, based on joint work with Carl-Christian Groh, I revisit a classic problem in mechanism design: the assignment of goods to agents. It is a well-established result that both the ex-post Pareto efficient and utilitarian optimal assignment rules assign the goods to the agents with the highest willingnesses to pay. In recent times, drastic inequality has become more and more prevalent and is being tack-

led by a variety of governments and organizations. To derive its implications on optimal mechanisms, I consider inequality in the classic mechanism design framework. Intuitively, if the utility of money exhibits decreasing returns, richer agents value money marginally less than poorer agents. I then show that it is necessary to account for inequality in utilitarian optimal mechanisms derive the optimal mechanism for the problem of goods assignment.

In particular, I study optimal mechanisms for a utilitarian designer who seeks to assign multiple units of an indivisible good to a group of agents. The agents have heterogeneous marginal utilities of money, which may naturally arise in environments where agents have different wealth levels or financing conditions. The designer faces constraints on ex ante transfers. I show that the ex post efficient allocation rule is not utilitarian optimal in the setting. In certain situations, it is utilitarian optimal to deterministically assign the good to an agent with a lower willingness to pay. This is because a high willingness to pay may stem from a low marginal utility of money. Moreover, the transfer rule does not only facilitate implementation of the desired social choice function in our setting, but also directly affects social welfare. Finally, I highlight how the mechanism can be implemented as an auction with minimum bids and bidding subsidies.

Partially Verifiable Information: Typically, there are at least two types of information that are of interest to a seller. First, how much a buyer is willing to pay for a good. Usually, this information is privately known by the buyer, and thus the buyer will enjoy information rents when buying from the seller. Second, auxiliary information which does not directly reveal the buyer's willingness to pay, but is informative. For example, where the buyer lives or how wealthy the buyer is. Another example is procurement, where a supplier's costs may depend on the type of machine with which a good is produced. Nowadays, an ever-increasing amount of data is produced, collected and used. Therefore, in the third chapter, I focus on the implications of auxiliary information, what I refer to as characteristics, on the optimal mechanism for the sale of a good.

In the model, I consider a seller selling a good to bidders with two-dimensional private information: their valuation for a good and their characteristic. While valuations are non-verifiable, characteristics are partially verifiable and convey information about the distribution of a bidder's valuation. I derive the revenue-maximizing mechanism and show that it can be implemented by introducing a communication stage before an auction. I show that granting bidders a right to remain anonymous, i.e., to refuse participation in the communication stage, leaves the optimal mechanism unchanged and provides no benefits for the bidders.

Acknowledgements

This dissertation contains my journey from a student to a researcher that I have undertaken in my time at the University of Mannheim. As with most journeys in life, it is exceedingly hard to reach a destination without any guidance and company along the way. Fortunately, I have enjoyed both. For their guidance, I want to thank my supervisor Thomas Troger, as well as my advisors Vitali Gretschko, Volker Nocke and Scott Duke Kominers.

During this journey, Thomas' advice has been extremely insightful, deep and meticulous. He always has an excellent grasp on which topics offer interesting research questions. In particular, I want to thank Thomas for getting me interested in Blockchains and Cryptocurrencies, which lead me to develop the idea for the first chapter of this dissertation. I want to thank Vitali not only for his excellent feedback, but also for welcoming me to the ZEW Market Design Group that has allowed me to broaden my horizon significantly. I want to thank Volker for patiently listening to my research at all stages, and for his incredible ability to grasp an unfamiliar model instantly and point me in the right direction. Finally, I want to thank Scott for being incredibly supportive and welcoming to a student who emailed him out of the blue asking for feedback.

The second chapter of this dissertation is based on joint work with my friend Carl-Christian Groh. It is the first full research paper that I have written. Working together has not only made the process much more enjoyable, but also taught me numerous valuable lessons that I am sure I would not have learned without him. I have also been fortunate to enjoy the company of many good friends in Mannheim - especially my fellow PhD students. Embarking on this journey together has allowed us to weather the storm together to arrive at our final destination. Thanks Christian, Federico, Jacopo, Jasmina, Li, Lukas (though technically, his journey had a different destination) and Mykola.

Finally, I want to thank my wife Tróc for her unwavering support throughout this journey. There are few certainties in life. But one thing that I know for certain is that I would not have reached my destination without her.

Chapter 1

The Value of Decentralization Using the Blockchain¹

Abstract: The popularity of blockchain technology and cryptocurrencies has grown in recent years, but there is still disagreement about their value in economic interactions. In this paper, I examine the value of a blockchain for an entrepreneur who creates a network. The entrepreneur can decide to retain control of the network with a centralized implementation through a regular company, or surrender control over the network with a decentralized implementation through the blockchain. The network's users experience a locked-in effect. I show that a decentralized implementation of the network is both (i) preferred by the entrepreneur and (ii) a Pareto improvement, if and only if the size of the locked-in effect is sufficiently large.

Keywords: Blockchain, Smart Contracts, Decentralization, Cryptocurrency, Commitment, Networks

JEL Classification: C70, D00, D2, D4, L2

1.1 Introduction

AWS (Amazon Web Services), Google, Facebook, Spotify, and Twitter are some of the largest tech companies that have billions of users worldwide. For an entrepreneur looking to create a competitor to these companies, the question arises: should they start a traditional company, or should they follow the path of decentralized networks like

¹I thank Thomas Troger for his continued support. I also thank Piotr Dworzak, Vitali Gretschko, Carl-Christian Groh, Federico Innocenti, Scott Duke Kominers, Volker Nocke, Marion Ott, Jonas von Wangenheim and audiences at the 2022 DICE Winter School for Applied Micro Theory, the 2022 CRC TR 224 Young Researchers Workshop and the 10th CRC TR 224 Retreat, NYU Stern, a16z and the EWMES for insightful comments. This work was supported by the University of Mannheim's Graduate School of Economic and Social Sciences. Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (Project B01) is gratefully acknowledged.

Filecoin, Internet Computer, Presearch, Decentralized Social, Minds, and Audius and leverage blockchain technology?² What is the benefit of decentralization and what incentives do entrepreneurs have to decentralize? Vitalik Buterin, co-founder of the Ethereum blockchain, argues that decentralization is, among other things, useful for Collusion resistance – it is much harder for participants in decentralized systems to collude to act in ways that benefit them at the expense of other participants, whereas the leaderships of corporations and governments collude in ways that benefit themselves but harm less well-coordinated citizens, customers, employees, and the general public all the time.³ Similar sentiments are shared throughout the white papers of several of the networks listed above.

The contribution of this paper is to develop a theoretical model that determines when an entrepreneur should implement a network in a centralized manner and when it is optimal to decentralize through the use of a blockchain. With that, I provide an answer to a question that is frequently raised when it comes to the topic of blockchain and cryptocurrencies: *Why should anybody use it?* As the core friction at play, I assume that users of the network are subject to a locked-in effect, for example due to switching costs.⁴ If the frictions that arise due to the potential of exploiting this locked-in effect by the entrepreneur are sufficiently large, I show that an entrepreneur prefers decentralizing her network. As a result, she effectively gives up control over the network and thus generates commitment to not abuse the locked-in effect of the users.

Achieving such commitment can lead to a Pareto improvement compared to a centralized implementation of a network through a regular company. That is, both the entrepreneur who creates the network, and the users may be better off if the network is decentralized. However, decentralization also comes at a cost for the entrepreneur: she surrenders the control over the network to the users and, to align incentives, engages in revenue sharing. Therefore, there is a trade-off between the costs of centralization and decentralization. I show that if the locked-in effect is small, an entrepreneur should implement her network in a centralized manner. On the other hand, if the locked-in effect is sufficiently large, an entrepreneur should implement her network in a decentralized

²<https://www.filecoin.io/> is "a decentralized storage network designed to store humanity's most important information". <https://www.internetcomputer.org/> is a "public blockchain that hosts smart contracts [...]". <https://www.presearch.io/> is a "decentralized search engine". <https://www.deso.org/> is "the decentralized social blockchain". <https://www.minds.com/> is an "open source social network dedicated to Internet freedom". <https://audius.org/> is a "decentralized protocol for audio content". Other examples of centralized companies with decentralized counterparts include: 1) payment processors such as Visa and various cryptocurrencies 2) centralized finance providers such as Banks and decentralized finance (DeFi) applications and centralized exchanges such as Binance and Coinbase and decentralized exchanges (Dex) such as Uniswap, Pancakeswap and others.

³<https://medium.com/@VitalikButerin/the-meaning-of-decentralization-a0c92b76a274>

⁴For example, Shapiro and Varian (1998) remark that switching costs are the norm, not the exception, in the information economy. For empirical measurements of switching costs, see for example Chen and Hitt (2002), Li and Agarwal (2017)

manner. Given this result, the list of companies with decentralized counterparts is not surprising. Arguably, users of AWS, Facebook, and other tech companies are subject to particularly large locked-in effects.

In the model, an entrepreneur (she) creates a network for her (potential) users (he). The users need the network to interact or achieve a goal. However, they lack the ability to develop a technological solution that suits their needs. The entrepreneur, on the other hand, possesses the necessary skills to build a network that fits the users' needs. At the start of the game, the entrepreneur decides between a centralized implementation of the network through a regular company and a decentralized implementation using the blockchain.⁵ In both implementations, the network can be monetized (for example through advertisement, sale of user data, or other means), and any revenues that are raised can be shared between the entrepreneur and the users. The entrepreneur and the users interact with each other through the network over an infinite time horizon. If the entrepreneur chooses centralized governance, she can change monetization and revenue sharing in every period. Each period, the existing users of the network have the choice to stay in the network or leave the network. Further, new users arrive every period and can choose to join the network. If the entrepreneur chooses decentralized governance, revenue sharing is decided by the entrepreneur through the tokenomics at the start of the game.⁶ Then, the users decide on monetization in every period through decentralized governance.⁷ As in centralized governance, each period, the existing users of the network have the choice to stay in the network or leave the network, and new users arrive who have the choice to join the network.

There is complete information, and the full history of the game is observed by both the entrepreneur and the users. The entrepreneur is purely interested in generating revenue through monetization, while the users' utility consists of three parts: First, they derive utility from using the network. Second, they dislike monetization such as advertisements, and third, they benefit from any revenue that is shared with them. I use sub-game perfect equilibria to analyze the game. Therefore, an entrepreneur using a centralized implementation of the network is unable to credibly commit to future levels of monetization and revenue sharing. Instead, her choice of monetization and revenue sharing has to be sequentially optimal for every history of the game given the strategy of the users.

⁵In an extension in section 1.3.1, I allow the entrepreneur to delay decentralization. That is, she can decentralize at a later time using an *airdrop*.

⁶Tokenomics is a mix of the two words token and economics. *Token* refers to a digital asset. The tokenomics then describe the underlying economics of that particular token, such as supply, distribution, vesting and other parameters.

⁷In practice, there are many mechanisms for on-chain governance. In the model, I use majority voting, where 1 unit of the token equals 1 vote, and an even split of tokens among the users.

I divide the analysis of the model into three subsections. First, the sub-game of centralized governance. Second, the sub-game of decentralized governance and third, determining the optimal governance structure for the network.

In the analysis of centralized governance, I show that the equilibrium of the game features two distinct phases. A growth phase in which new users join the network, and an exploitation phase in which no new users join the network and the entrepreneur exploits the locked-in effect of the existing users through increased monetization and decreased revenue sharing. The transition between the two phases crucially depends on network effects and the network's future growth, and is characterized by the point at which the entrepreneur is indifferent between attracting new users and foregoing growth to exploit the locked-in effect of the existing users. In equilibrium, the users anticipate being locked-in to the network and have to be compensated up front to be incentivized to join the network.⁸ The compensation equals the discounted value of the switching costs that lead to the locked-in effect. Thus, as the severity of the locked-in effect increases, it becomes increasingly harder for the entrepreneur to attract users in the first place. I show that for a sufficiently large locked-in effect, no users join the network in equilibrium, resulting in zero revenues for the entrepreneur. This highlights the commitment problem, that an entrepreneur may try to solve with decentralization through a blockchain.⁹

If the entrepreneur chooses decentralized governance, the degree of monetization is decided by the users. Unlike the entrepreneur, the users internalize the negative effects of monetization through their utility function. As a result, the locked-in effect will not be exploited when the monetization of the network is controlled by the users. To align incentives, the entrepreneur engages in revenue sharing with the users. Further, the network grows every period, unlike in centralized governance. However, decentralized governance has two drawbacks. First, the entrepreneur surrenders control over the network, such that she cannot choose the degree of monetization she prefers. Second, because users choose the degree of monetization, the entrepreneur has to engage in revenue sharing to align incentives.

Finally, I determine the optimal governance structure of the network by comparing centralized governance to decentralized governance. I show that for minimal locked-in effects, an entrepreneur is better off choosing centralized governance. In contrast, for a sufficiently large locked-in effect, decentralized governance is preferred, as the entrepreneur

⁸This property of the equilibrium is nicely summarized in Shapiro and Varian (1998)'s advice to buyers that anticipate becoming locked-in: Bargain hard at the outset of the lock-in cycle for a sweetener or some form of long-term protection before you become locked in

⁹An alternative solution to creating commitment for a centralized network could be contracting over monetization and revenue sharing. However, it is likely that these contracts would be incomplete. Thus, contracting may face issues such as renegotiation, as discussed in the literature on incomplete contracts (e.g., Hart and Moore (1988), Hart and Moore (1999)), and fail to be a suitable solution.

is unable to attract any users when choosing centralized governance. To determine the optimal mode of governance for an arbitrary size of the locked-in effect, I show that the revenue that the entrepreneur can achieve with centralized governance is a decreasing function of the size of the locked-in effect. In contrast, the revenue that the entrepreneur can raise with decentralized governance is independent of the size of the locked-in effect. Thus, there exists a threshold size, such that the entrepreneur should decentralize her network if and only if the locked-in effect is sufficiently severe.

Literature: This paper contributes to the literature on the economics of blockchains. It most closely relates to papers that have discussed blockchain technology regarding commitment and competition. Similar to Sockin and Xiong (forthcoming), I consider an entrepreneur who can exploit the platform's users and show that creating commitment through the blockchain may be beneficial for the entrepreneur. My paper contributes relative to theirs as follows: First, they consider a one shot interaction between the entrepreneur and the users on the platform. As such, in centralized governance, exploitation occurs for sure since there is no ongoing relationship between the entrepreneur and the users. I contribute by considering a repeated interaction between the entrepreneur and the users, and show that the problem of exploitation persists even in repeated interactions. Further, I consider the potential for user growth in the network, and show that user growth can be a substitute for commitment when future growth is strong, but fails to generate commitment when future growth is sufficiently low. Finally, the longer time horizon allows me to consider locked-in effects and show that the entrepreneur decentralizes her network if and only if the locked-in effect is sufficiently large.

Goldstein et al. (2019) argue that using an initial coin offering (ICO) and committing to the free resale of tokens can enable a monopolistic entrepreneur to commit to competitive pricing. My paper complements their contribution by focussing on the importance of locked-in effects in platforms. Both papers demonstrate that commitment through the blockchain may improve welfare. However, Goldstein et al. (2019) show that committing to the free resale of tokens yields lower profits for an entrepreneur compared to operating the network in a traditional, centralized manner. In contrast, I show that an entrepreneur can increase her revenue by implementing her network through the blockchain, if the costs of centralization are too large. Further, I contribute by adding network growth and showing that growth can be a substitute for commitment at first, but fails to be a substitute for commitment when growth slows down over time.

Huberman et al. (2021) focus on bitcoin as a payment system (BPS), which can be considered as a network in the terms of my model, and show that user surplus in the BPS is larger compared to a monopolist payment provider. However, the incentives for a

monopolist to set up a decentralized network such as bitcoin remain unclear. Brzustowski et al. (2021) show that the Coase conjecture fails if a seller can generate commitment through smart contracts.

Catalini and Gans (2018) focus on entrepreneurs that are capital constrained and need to raise capital through an ICO to fund their network. Bakos and Halaburda (2018), Li and Mann (2018) and Cong et al. (2021), show how ICOs can mitigate coordination failures in the users' decision to join or not join a particular network. In empirical assessments of ICOs, Howell et al. (2020) find that success in ICOs is associated with disclosure, credible commitment to the network, and quality signals, while Adhami et al. (2018) find that, among other things, revenue sharing makes ICOs more successful.

Arrunada and Garicano (2018) and Chen et al. (2021) investigate the details of decentralized governance more closely. Further, this paper also relates to the literature of blockchain consensus, as it shares some intersections with blockchain governance. Contributions include Abadi and Brunnermeier (2018), Biais et al. (2019), Catalini et al. (2020) and Saleh (2021). Decentralization through the blockchain gives users decision power in the network. Thus, my paper also shares some commonalities with the literature on common ownership in traditional corporations, for example Magill et al. (2015), Cres et al. (2020) and Azar and Vives (2021), but with a drastically different focus.

Another strand of the literature that connects to my model is the IO literature on (two-sided) platforms and network effects, as all the applications I have mentioned are platforms, with seminal contributions by Katz and Shapiro (1985), Farrell and Saloner (1986), Rochet and Tirole (2003) and Armstrong (2006). Cabral (2011) develops a dynamic model of platform competition.¹⁰ This literature focuses on equilibrium pricing and competition between platforms. As such, my paper is complementary, as my model features neither competition between platforms nor focuses on prices for either side of the market. I focus on the value of commitment for the entrepreneur as a function of the size of the locked-in effect of the platform. I also connect to papers that - from a regulatory perspective - investigate platform governance, for example Jullien and Pavan (2019), Choi and Jeon (2022) and Teh (2022). For a general overview of the literature, see for example Farrell and Klemperer (2007) and Belleflamme and Peitz (2021).

The rest of the paper is structured as follows: Section 1.2 consists of the model and the results that outline when decentralization through the blockchain is preferable to centralization. Section 1.3 discusses extensions of the model. Section 1.4 concludes. For readers that are not familiar with blockchains, appendix 1.5.1 provides a supplementary overview over the blockchains, some use cases, and how it enables an entrepreneur to

¹⁰Peitz et al. (2017) study price setting dynamics on platforms experimentally.

generate commitment.

1.2 Model

The model is a sequential game with infinitely many periods $t = 0, 1, 2, \dots$ between an entrepreneur (she) and a continuum of users (he), indexed by i . The entrepreneur creates a network for the users in $t = 0$ and the mass of users in the network at time t is denoted by μ_t . In $t = 1, 2, 3, \dots$ the network can be monetized (for example through advertisement, sale of user data, or other means). The revenue from monetization can be decomposed into two parts. First, there is a level of monetization of the network $\pi_t \in \mathbb{R}_+$. This variable represents the intensity with which the network is monetized, such as how often or how many advertisements are displayed, or how much of the user data is sold. Second, given a measure of users μ_t and a level of monetization π_t , the revenue generated by the network equals $\pi_t \phi(\mu_t)$ where ϕ is an increasing, continuously differentiable function with $\phi(0) = 0$. $\phi(\mu_t)$ represents the rate an advertiser is willing to pay for advertisements or for user data. Throughout the paper, I assume that $\frac{\phi(\mu_t)}{\mu_t}$ is non-decreasing in μ_t .¹¹ Any revenues that are raised can be shared between the entrepreneur and the users. The fraction of revenue that the entrepreneur keeps is denoted by α_t , while the leftover fraction of revenue $(1 - \alpha_t)$ is shared with the users.

How monetization and revenue sharing are chosen depends on the mode of governance of the network. At the beginning of the game, in $t = 0$, the entrepreneur chooses the mode of governance (centralized or decentralized). If the entrepreneur chooses centralized governance, she can change monetization π_t and revenue sharing α_t in every period $t = 1, 2, \dots$. Each period, users have a binary choice. The existing users of the network have the choice to stay in the network or leave the network. Further, new users arrive every period and can choose to join or not join the network.

If the entrepreneur chooses decentralized governance, she commits, without loss of generality, to a fixed percentage α of revenue sharing in $t = 0$ through the tokenomics of the network.¹² She achieves this through the appropriate distribution of the network's token between herself and the users.¹³ In every period $t = 1, 2, \dots$ the users of the network determine the amount of monetization π_t through *on-chain governance*. As in centralized governance, each period, users have a binary choice. The existing users of the network

¹¹For example, this holds true in cost-per-view and cost-per-click advertisement that is commonly used in online advertisement. If c is the cost per click/view and a fraction $\gamma \in [0, 1]$ of the users interacts with advertisement, it holds that $\frac{\phi(\mu_t)}{\mu_t} = c\gamma$, which is constant in μ_t .

¹²In an extension in appendix 1.5.12, I allow the entrepreneur to pre-commit to a path for revenue sharing and show that she chooses a constant percentage of revenue sharing. Thus, considering a fixed percentage throughout the main body of the paper is without loss of generality.

¹³For the example of Uniswap, 60% of the token supply has been allocated to users, while the other 40% is split between the Uniswap team, investors, and advisors. For details, see <https://uniswap.org/blog/uni>.

have the choice to stay in the network or leave the network. Regardless of the mode of governance, users that decide to leave the network or newly arriving users who decide not to join the network drop out of the game and realize the value of their outside option.

Every period, new potential users become aware of the network. Let μ_{t-1} be the mass of users in the network in period $t - 1$. Then, in period t there will be a mass of $g(\mu_{t-1}) - \mu_{t-1} \geq 0$ new users who become aware of the network. Each potential new user has the choice to join or not join the network. For example, if all new users join, the new measure of users in the network is equal to $g(\mu_{t-1})$. If no new user joins, the network remains at μ_{t-1} users. The growth function g is continuously differentiable and the mass of users in period 0 is equal to $\mu_0 = 0$. If the network loses all its users within a period, no new users will arrive at any point in the future. This assumption rules out cyclical equilibria in which the entrepreneur continuously "starts over". There is complete information and both the entrepreneur and the users observe the full history of the game.

The entrepreneur is strictly interested in revenue: her utility in a particular period t is equal to her revenue share α_t multiplied by the revenue raised by monetization $\pi_t \phi(\mu_t)$: $u_t^E = \alpha_t \pi_t \phi(\mu_t)$. The utility a user receives from participating in the network has three components: First, a user derives utility $V(\mu_t)$ from using the network. I assume that V is increasing, continuously differentiable and that $V(0) = 0$. Second, as a result of the monetization of the network, π_t , the user's utility decreases by $k\pi_t^2$, where $k > 0$ describes the user's aversion to monetization. This represents the decrease in utility a user suffers when being forced to watch advertisements, through the sale of his data, or other detrimental effects of monetization. As a third component, a user may potentially receive a share of the revenues that the network generates. I assume that this share is equally split between all users, such that each user receives a fraction $\frac{1-\alpha_t}{\mu_t}$ of the revenue. The utility function of a user thus equals $u_t = V(\mu_t) - k\pi_t^2 + \frac{(1-\alpha_t)}{\mu_t} \pi_t \phi(\mu_t)$.

A user who newly arrives in the network can decide to join the network and realize the utility as described above. If the user decides not to join the network, he realizes an outside option that is normalized to 0. A user who has already taken part in the network for at least one period can decide to stay in the network, realizing the utility of participating, or leave the network. However, the outside option for these users is equal to $-u < 0$. Thus, users that already take part in the network suffer from a *locked-in effect*. This assumption represents the idea that users have spent time interacting with the network, such that its algorithm has adapted to their needs.¹⁴ An equivalent interpretation is that

¹⁴For example, Google's search algorithm learns from a user's past searches and improves its search results. Spotify's algorithm learns a user's taste in music, improving the likelihood of playing music that the user likes.

the value of the outside option has remained constant, but users encounter a switching cost equal to u when leaving the network in favor of the outside option.

Both the entrepreneur and the users maximize the sum of their discounted utilities. Future utilities are discounted by a common discount factor $\delta \in (0, 1)$. I divide the analysis into subsections dedicated to the sub-games of centralized and decentralized governance. Within those sections, I give a detailed description of the structure of the sub-games of centralized and decentralized governance. Then I derive the sub-game perfect Nash equilibria and discuss their properties. Finally, I determine the optimal decision of the entrepreneur at the start of the game: to choose centralized or decentralized governance for her network.

1.2.1 Centralized Governance

If the entrepreneur chooses centralized governance, every period $t = 1, 2, \dots$ has the following timing:

1. The entrepreneur chooses a level of monetization π_t and a fraction of revenue sharing α_t
2. Users make a simultaneous choice:
 - (a) Users that arrived in period t choose to join or not to join
 - (b) Users who are already present in the network choose to stay or leave
3. Utilities realize

A centralized entrepreneur retains full control over the monetization and revenue sharing of the network. However, she lacks the ability to commit to the levels of monetization and revenue sharing for future periods because her strategy has to be sequentially optimal. Now I can define strategies for the entrepreneur and the users in more detail. For that, define by h_t a history of the game up to period t . Then a strategy is defined as a mapping from the set of possible histories into the possible actions. Specifically, for the entrepreneur, a strategy maps any possible history into some degree of monetization π_t and revenue sharing α_t . For the users, a strategy maps into the binary decisions to join or to not join at their time of arrival in the network, or, if already present in the network, into a binary decision of staying or leaving. I impose the following tie-breaking rules: Newly arriving users that are indifferent between two strategies, such that one prescribes joining the network and one prescribes not joining the network will join the network. Users that are indifferent between two strategies, such that one strategy prescribes leaving the network and another strategy prescribes not leaving in the network, will choose

to remain in the network.

As a preliminary step in the analysis, it is useful to think about optimal choices of monetization and revenue sharing within a given period. That is, what choice of monetization and revenue sharing maximizes the entrepreneur's revenue, given that the users should receive some arbitrary level of utility \hat{u} , and how large is the corresponding revenue for the entrepreneur. The result is derived from a standard constrained optimization problem. From now on, I will denote the entrepreneur's revenue that results from the optimal choice of monetization and revenue sharing for a network of size μ_t with user utility level \hat{u} by $\psi(\mu_t, \hat{u})$. This function $\psi(\mu_t, \hat{u})$ will be crucial for the analysis of centralized governance. For brevity, the derivation of $\psi(\mu_t, \hat{u})$ is relegated to appendix 1.5.2. In the main body of the paper, I focus on describing the characteristics of $\psi(\mu_t, \hat{u})$ and providing some intuitions. First, the entrepreneur's revenue is increasing in the amount of users μ_t and decreasing in the level of utility \hat{u} that the users receive. As such, there is a conflict of interest between the entrepreneur and the users. Second, there is a limit to how large the user utility level \hat{u} can be for a given network size μ_t . It is not feasible to provide a user utility level that exceeds what a user would receive if the entrepreneur distributed the entire revenue to the users. Last, depending on the users' aversion to monetization k , the centralized network may feature revenue sharing. That is, for small values of k , the entrepreneur will increase the monetization of the network and compensate the users by sharing some of the revenue. In contrast, when k is large, the entrepreneur will not share any revenue with the users.

To derive the equilibrium of the centralized governance sub-game, it is instructive to consider the entrepreneur's incentives to grow her network. Every period, new users arrive to join the network potentially. For the network to grow, joining the network has to be weakly beneficial for a newly arriving user. That is, joining the network has to yield at least utility equal to 0. Instead of growing the network, the entrepreneur can exploit the existing users. Given that existing users are locked into the network and have an outside option that is valued at $-u < 0$, the entrepreneur can potentially achieve a higher level of revenue when focusing on extracting additional revenue from existing users. To quantify the revenue that an entrepreneur generates when she decides to exploit the users in her network, consider some period t . The amount of existing users at the start of the period is equal to μ_{t-1} . If she exploits the existing users forever, the present value of the stream of her discounted future revenue equals

$$\frac{1}{1-\delta}\psi(\mu_{t-1}, -(1-\delta)u) \quad (1.2.1)$$

Note that the entrepreneur provides a per-period utility of $-(1-\delta)u$ to the users, such that

the discounted utility is equal to $-u$, keeping the users indifferent between staying and leaving. To grow the network, the entrepreneur has to provide enough utility to the users, such that they are better off joining the network in the first place. If the entrepreneur grows the network one last time in some period t before exploiting the existing users, she has to provide utility δu to the last users who are to join the network. The entrepreneur's revenue from growing the network one more time and then exploiting the network's users from that point onward equals

$$\psi(g(\mu_{t-1}), \delta u) + \frac{\delta}{1-\delta} \psi(g(\mu_{t-1}), -(1-\delta)u) \quad (1.2.2)$$

The point at which the entrepreneur is indifferent between growing the network one last time and exploiting the existing users in her network will be crucial for the analysis of the equilibrium. I denote the solution to the following equation by $\bar{\mu}$:

$$\frac{1}{1-\delta} \psi(\bar{\mu}, -(1-\delta)u) = \psi(g(\bar{\mu}), \delta u) + \frac{\delta}{1-\delta} \psi(g(\bar{\mu}), -(1-\delta)u) \quad (1.2.3)$$

It is exactly at the network size $\bar{\mu}$ where the entrepreneur is indifferent between growing the network one last time and then exploiting the users in the future, and exploiting the users right away. It highlights the trade-off between exploiting the locked-in effect of a smaller mass μ_{t-1} of users starting today, or, growing the network at the cost of providing utility δu to the users to then exploit a larger network with $g(\mu_{t-1})$ users starting tomorrow. For the purpose of this paper, I focus on the case where such a value $\bar{\mu}$ exists. Indeed, this captures the economically interesting case of the model. If no such $\bar{\mu}$ exists, the entrepreneur never wants to exploit her users, regardless of how many users there are to exploit and how few users will arrive in the future. In appendix 1.5.3 I provide an extensive discussion of sufficient conditions to assure that $\bar{\mu}$ is well-defined. For the main body of the paper, I focus on providing an intuitive characterization of these settings. The key feature is the idea, that user growth will slow down over time. Indeed, if the overall pool of potential users is limited and a large amount of users has already joined the network, user growth necessarily slows down mechanically over time. However, there is some nuance in that a slowdown in user growth can be partially offset through an increase in revenues due to network effects. If these network effects are particularly strong relative to the growth rate of the network, growing the network remains preferable for the entrepreneur. What is important for $\bar{\mu}$ to exist, is that eventually growth slows down sufficiently to offset increased network effects, or that the network effect of attracting an additional eventually diminishes when the network is large. As a last point, I want to provide one particularly tractable example: $V(\mu_t)$ is constant, $\phi(\mu_t)$ is linear in μ_t and $g(\mu_t) = \mu_t + \gamma(\mu_t)$ where $\gamma(\mu_t)$ is a strictly decreasing, strictly positive function that approaches 0 as $\mu_t \rightarrow \infty$

For the following analysis, suppose that

$$\frac{\phi(\bar{\mu})^2}{4k\bar{\mu}^2} + V(\bar{\mu}) \geq \delta u \quad (1.2.4)$$

This condition ensures that it is feasible to the entrepreneur to ensure the utility level δu to a network of size $\bar{\mu}$. Later, I discuss what happens when this condition is not satisfied. For a better understanding of the equilibrium that will follow shortly, I want to emphasize that the level of user utility \hat{u}_t that is implied by a degree of monetization π_t and revenue sharing α_t is a function of the amount of users μ_t that are present in the network at the end of period t . For example, a particular tuple (π_t, α_t) implies different user utility levels \hat{u}_t when $\mu_t = 0$ compared to when $\mu_t > 0$. Now, the intuition of the trade-off between growing the network and exploiting the existing users can be condensed into an equilibrium:

Proposition 1 *Suppose condition 1.2.4 is satisfied. Then the following strategies constitute a sub-game perfect Nash equilibrium:*

Entrepreneur's strategy:

- If $\mu_{t-1} < g^{-1}(\bar{\mu})$, set π_t and α_t to maximize revenue as given by $\psi(\mu_t, \hat{u}_t)$ for user utility level $\hat{u}_t = 0$ and network size $\mu_t = g(\mu_{t-1})$
- If $g^{-1}(\bar{\mu}) \leq \mu_{t-1} < \bar{\mu}$, set π_t and α_t to maximize revenue as given by $\psi(\mu_t, \hat{u}_t)$ for user utility level $\hat{u}_t = \delta u$ and network size $\mu_t = g(\mu_{t-1})$
- If $\bar{\mu} \leq \mu_{t-1}$ set π_t and α_t to maximize revenue as given by $\psi(\mu_t, \hat{u}_t)$ for user utility level $\hat{u}_t = -(1 - \delta)u$ and network size $\mu_t = \mu_{t-1}$

Users' strategy:

- In the period of arrival, join the network iff
 1. $\mu_{t-1} < g^{-1}(\bar{\mu})$ and π_t, α_t are such that user utility level $\hat{u}_t \geq 0$ for a network size $\mu_t = g(\mu_{t-1})$
 2. $g^{-1}(\bar{\mu}) \leq \mu_{t-1}$ and π_t, α_t are such that user utility level $\hat{u}_t \geq \delta u$ for a network size $\mu_t = g(\mu_{t-1})$
- If already locked in to the network, stay in the network iff π_t, α_t are such that user utility level $\hat{u}_t \geq -(1 - \delta)u$ for a network size $\mu_t \geq \mu_{t-1}$

Proof. See appendix 1.5.4 ■

The equilibrium features the cutoff $\bar{\mu}$, at which the entrepreneur switches from growing the network to exploiting the existing users in the network. The entrepreneur's strategy has three distinct parts. If $\mu_{t-1} < g^{-1}(\bar{\mu})$, the entrepreneur will grow the network again in the next period, as $g(\mu_{t-1}) < \bar{\mu}$. Thus, the entrepreneur sets user utility equal to $\hat{u}_t = 0$ and the users are willing to join the network. Note that in these periods, the entrepreneur has basically regained commitment to not abuse the locked-in effect of the users. The entrepreneur refrains from exploiting the locked-in effect of the existing users in the network with the aim to grow the network larger. At $g^{-1}(\bar{\mu}) \leq \mu_{t-1} < \bar{\mu}$, the entrepreneur reaches the limits of how far she is willing to grow the network. If the entrepreneur grows the network it holds that $\mu_t = g(\mu_{t-1}) > \bar{\mu}$, such that in the future, the entrepreneur will be better off with exploiting the locked-in effect of the users compared to growing the network any further. However, to attract users to the network, the entrepreneur has to offer a utility level equal to $\hat{u}_t = \delta u$. In the last part, when $\bar{\mu} \leq \mu_{t-1}$, the entrepreneur is better off exploiting the locked-in effect of the network's existing users compared to growing the network any further.

The users' strategies are as follows: when they newly arrive at the network, they do not suffer from a locked-in effect. They observe the network size and if $\mu_{t-1} < g^{-1}(\bar{\mu})$, anticipate that the entrepreneur will grow the network further in the future, such that it is optimal for them to join the network if $\hat{u}_t \geq 0$. If $g^{-1}(\bar{\mu}) \leq \mu_{t-1}$, they know that the entrepreneur will grow the network at most one more time. As such, they require a level of utility at least equal to δu to join the network. If they are already locked into the network, they will remain in the network iff $\hat{u}_t \geq -(1 - \delta)u$, as this implies that the discounted value of their future utility is at least equal to the value of their outside option $-u$. Note that no profitable deviations exists for neither the entrepreneur nor the users. In equilibrium, newly arriving users are indifferent between joining and not joining the network, while users that are already locked into the network strictly prefer staying in the network before the entrepreneur starts exploiting the users and are indifferent between staying and leaving when the entrepreneur starts exploiting the network. For the entrepreneur, deviations that increase the users' utility level are not profitable, since it does not change the users actions on the equilibrium path and her revenues are decreasing in the users' utility levels. Decreasing the utility offered to the users at any point in time will cause the users to leave the network, resulting in 0 revenues, thus not being a profitable deviation.

Now reconsider what happens if

$$\frac{\phi(\bar{\mu})^2}{4k\bar{\mu}^2} + V(\bar{\mu}) < \delta u \quad (1.2.5)$$

Then, the entrepreneur cannot pay the compensation utility δu in the last period where she will grow the network. If the entrepreneur sets a utility level of less than δu , no new users will join, as the value of joining is below the outside option of 0. However, if the entrepreneur is unable to attract any new users, she should maximize revenues from the existing users of the network. That is, setting user utility equal to $-(1 - \delta)u$ instead. Denote this last period of potential growth in which this issue occurs as t^* . Then, users should anticipate that the entrepreneur will exploit the locked-in effects not starting from period $t^* + 1$ onward, but from period t^* . Then, the users who arrive at period $t^* - 1$ need to be provided utility level δu , for them to be incentivized to join the network. However, note that at period $t^* - 1$ the size of the network is necessarily smaller than at t^* . Thus, since the network's revenues are increasing in the mass of user μ_t , it is also not feasible for the entrepreneur to provide utility level δu to the users in period $t^* - 1$. This logic carries forward until the first period, such that no users should join the network at all. To further examine when this issue occurs, define by $\underline{\mu}$ the solution to the equation

$$\frac{\phi(\underline{\mu})^2}{4k\underline{\mu}^2} + V(\underline{\mu}) = \delta u \quad (1.2.6)$$

Intuitively speaking, $\underline{\mu}$ is the minimum required size of the network, such that it is feasible for the entrepreneur to provide utility δu to the users. Now, if $\bar{\mu} \geq \underline{\mu}$, the case discussed above does not occur and the entrepreneur can attract users to her network. However, if $\bar{\mu} < \underline{\mu}$, the entrepreneur is unable to attract any users to her network. The entrepreneur's main issue in the network with centralized governance is her lack of commitment to not abusing the locked-in effect of the users. Thus, I focus on the effects of the severity of the locked-in effect u on $\underline{\mu}$ and $\bar{\mu}$.

Lemma 1 $\underline{\mu}$ strictly increases in u . As $u \rightarrow \infty$ it holds that $\underline{\mu} \rightarrow \infty$.

To see why the lemma holds true, consider equation 1.2.6. When u increases, the RHS of the equation increases. Then the lemma clearly holds true, as the LHS of the equation is increasing in $\underline{\mu}$ since $\frac{\phi(\underline{\mu})^2}{4k\underline{\mu}^2}$ is increasing in $\underline{\mu}$ (recall that $\frac{\phi(\underline{\mu})}{\underline{\mu}}$ is increasing in $\underline{\mu}$ by assumption) and $V(\underline{\mu})$ is also increasing in $\underline{\mu}$ by assumption.

Next, consider $\bar{\mu}$. Note that $\bar{\mu}$ is only implicitly defined in equation 1.2.3. It is the size of the network that makes the entrepreneur indifferent between growing the network once more today and exploiting the users in the future vs. exploiting the users starting today. As such, I employ the implicit function theorem to show the following lemma:

Lemma 2 $\bar{\mu}$ strictly decreases in u . As $u \rightarrow 0$ it holds that $\bar{\mu} \rightarrow \infty$.

Proof. See appendix 1.5.5. ■

As the size of the locked-in effect grows, the entrepreneur stops growing the network and start exploiting the existing users earlier. With a larger locked-in effect, there is

more to gain by exploiting the existing users. To sum things up, I have shown that $\underline{\mu}$ is strictly increasing in u and that $\bar{\mu}$ is strictly decreasing in u . Therefore, as u increases, the following two effects take place. First, the entrepreneur needs a larger size network to make it feasible to guarantee users a utility level δu in the last period of growth. Second, as u increases, the entrepreneur is more tempted to exploit the existing users of the network and stops growing the network earlier. Therefore, the following corollary formalizes that when u grows too large, the entrepreneur is unable to attract any users to her network:

Corollary 1 *There exists some value u^* such that the entrepreneur is unable to attract any users to the network if $u > u^*$. Consequently, the equilibrium revenue of the network with centralized governance is 0.*

The corollary follows by defining u^* as the value of u for which $\underline{\mu} = \bar{\mu}$. Then for all $u > u^*$ it holds that $\bar{\mu} < \underline{\mu}$. As the size of the locked-in effect grows too large, the entrepreneur will more readily exploit users who are already in the network, rather than growing the network by attracting new users. However, in equilibrium, this is anticipated by any users that arrive at the network, such that no users join the network at all. This highlights the commitment problem of the entrepreneur. If she was able to commit to not abusing the locked-in effect of the users, she would be able to attract users to her network and generate revenues. Note that this corollary establishes a sufficiency result. When the size of the locked-in effect is sufficiently large, it is better to decentralize the network, if the entrepreneur can attract at least some users in decentralized governance. In section 1.2.3 I show that this result carries over more generally, by determining the cutoff size for the locked-in effect such that the entrepreneur prefers to decentralize the network if and only if the locked-in effect is sufficiently severe. Before that, the next section discusses the sub-game of decentralized governance.

1.2.2 Decentralized Governance

If the entrepreneur chooses decentralized governance, every period $t = 1, 2, \dots$ has the following timing:

1. Users make a simultaneous choice:
 - (a) Users who are not present in the network choose to join or not to join
 - (b) Users who are already present in the network choose to stay or leave
2. Users collectively choose π_t
3. Utilities realize

This section focuses on the sub-game of decentralized governance. First, the entrepreneur chooses, without loss of generality, a permanent revenue split α . Then, users that have newly arrived have the choice to join or not join the network. Existing users have the choice to stay or leave the network. Afterward, users vote on the degree of monetization π_t for the period and utilities realize. When analyzing the voting equilibria, I will restrict the equilibrium analysis to weakly undominated strategies. In voting games, the strategy of voters has to be optimal, conditional on being pivotal. As no single voter is ever pivotal when there is a continuum of users, basically any strategy can be played in an equilibrium. Therefore, restricting the users' strategies to be weakly undominated, implies that they truthfully vote for their preferred degree of monetization π_t as if they were pivotal. This leads to the following equilibrium:

Proposition 2 *There is a sub-game perfect equilibrium such that every period the users of the network will vote for a degree of monetization*

$$\pi_t^* = \frac{1 - \alpha \phi(\mu_t)}{2k \mu_t} \quad (1.2.7)$$

The network will grow every period. The entrepreneur shares half of the revenue with the users.

Proof. See appendix 1.5.6. ■

The equilibrium highlights that decentralized governance is an effective commitment tool for the entrepreneur. In contrast to centralized governance, the users can be certain that their locked-in effect will not be exploited by the entrepreneur. Thus, users will continue to join the network every period. However, for the entrepreneur, this commitment comes at a substantial cost: she shares half the revenues of the network with her users. Nonetheless, it is necessary for her to share revenue with her users. If she would not share any revenue, the users would subsequently vote to stop the monetization of the network. As a result, the entrepreneur would not receive any revenue. Therefore, the sharing of revenue in a decentralized implementation of the network is necessary, as it aligns the incentives of the entrepreneur and the incentives of the network's users.

One potential point of contention in decentralized governance could be conflicts of interest between existing and newly arriving users. The users' utility function equals $V(\mu_t) - k\pi_t^2 + \frac{1-\alpha}{\mu_t} \pi_t \phi(\mu_t)$. The share of revenue that each user gets in the network is $\frac{1-\alpha}{\mu_t}$. As such, newly arriving users will dilute the revenue shares of existing users in the network. However, note that the users' per period utility in the equilibrium equals $V(\mu_t) + \frac{\phi(\mu_t)^2}{8k\mu_t^2}$. Since $\frac{\phi(\mu_t)}{\mu_t}$ is non-decreasing by assumption, the equilibrium utility is increasing in μ_t . Intuitively speaking, the network effects that accompany the entry of new users sufficiently compensate the dilution of the revenue share of existing users. Thus,

there is no incentive for existing users to try to prevent entry from newly arriving users to avoid dilution of their revenue shares.

1.2.3 Optimal Governance

The two preceding sections have solved the sub-games of centralized and decentralized governance. Now the main question remains: which form of governance the entrepreneur should choose when she creates her network? As has been shown in proposition 1, centralized governance will result in the entrepreneur eventually stopping to grow the network and starting to exploit the locked-in effect of the users. This change from network growth to exploiting the users is inherent in centralized governance, as the entrepreneur is unable to commit to future monetization and revenue sharing. Subsequently, corollary 1 showed that, when the locked-in effect is sufficiently large, the entrepreneur is unable to attract any users to the network, yielding her 0 revenue in equilibrium. This threshold of the locked-in effect serves as a sufficient condition for when it is optimal to decentralize. However, a complete comparison between the entrepreneur's revenue in centralized and decentralized governance remains. That is, what is the optimal mode of governance for any arbitrary size of the locked-in effect? To answer this question, I start by considering the opposite extreme of what was discussed in the corollary, namely when the locked-in effect is very small. Then, I move to locked-in effects of arbitrary size.

For small locked-in effects, the commitment problem of the entrepreneur becomes less and less severe, and in the limit of $u = 0$, disappears entirely. Comparing centralized and decentralized governance for $u = 0$ is rather straightforward. When $u = 0$, there is no locked-in effect that can be abused by the entrepreneur in the future. Thus, users will join the network every period, resulting in growth in any period in the centralized network. In comparison, note that the decentralized network also featured growth in every period. As such, the potential revenues that can be generated in both modes of governance are the same. However, in centralized governance, the entrepreneur stays in control and can generate maximum amounts of revenue for herself, while she surrenders control over the network in decentralized governance and has to engage in revenue sharing to align the users' preferences with hers. Thus, centralized governance is superior when the locked-in effect is small. This intuition is condensed in the following lemma:

Lemma 3 *As $u \rightarrow 0$ centralized governance is always preferred over decentralized governance.*

Proof. See appendix 1.5.7 ■

So far, I have established comparisons of centralized and decentralized governance at both extremes of the size of the locked-in effect. For minimal locked-in effects, centralized

governance is optimal for the entrepreneur, while for sufficiently large locked-in effects, decentralized governance is optimal for the entrepreneur. For intermediate values, the optimal mode of governance is hard to compute explicitly, as the revenue of the entrepreneur in the centralized network is only given implicitly, through the implicit definition of the maximum network size $\bar{\mu}$. However, what can be shown is a monotonicity result. That is, as the size of the locked-in effect increases, the entrepreneur's revenue in centralized governance decreases. As a result, there is a clear cutoff in the size of the locked-in effect, such that decentralized governance is preferred if and only if the size of the locked-in effect is larger than this cutoff. This idea is condensed into the following proposition:

Proposition 3 *There exists a well-defined size of the locked-in effect, u^{**} , such that decentralized governance is preferred by the entrepreneur if and only if $u > u^{**}$.*

Proof. See appendix 1.5.8. ■

The idea of the proof is as follows. First, recall that I have shown that at the two extremes of minimal and very large locked-in effects, the entrepreneur prefers centralized and decentralized governance respectively. Next, note that the entrepreneur's revenue with decentralized governance is independent of the size of the locked-in effect u . This holds as the users decide the level of monetization in the network with decentralized governance, and their optimal decision does not depend on u . The final step of the proof shows, that the entrepreneur's revenue with centralized governance is decreasing in the size of the locked-in effect u . Together, these observations imply the result, as they imply that the functions of the revenue under centralized and decentralized governance can cross at most once.

To realize why the entrepreneur's revenue with centralized governance is decreasing in u , consider the effect of a change in the size of the locked-in effect. In the centralized network, revenue is generated in three different phases. First, is the growth phase in which the entrepreneur provides 0 period utility to the users. Second, the last period of growth in which the entrepreneur provides utility equal to δu to the users, and lastly, the periods of exploiting where the entrepreneur provides utility equal to $-(1 - \delta)u$ to the users. Consider the immediate effect of an increase in u . The revenues of the first phase of the network are independent of u and remain unchanged. Second, the required period utility of the users in the last phase of growth, δu increases, resulting in decreased revenue for the entrepreneur. Finally, the user utility level in the exploitation phase, $-(1 - \delta)u$ decreases and leads to increased revenues for the entrepreneur. However, the entrepreneur's revenue is a function that is concave in the utility level (c.f. $\psi(\mu_t, \hat{u}) = \mu_t V(\mu_t) + \frac{\phi(\mu_t)^2}{4k\mu_t} - \mu_t \hat{u}$ or $\psi(\mu_t, \hat{u}) = \sqrt{\frac{V(\mu_t) - \hat{u}}{k}} \phi(\mu_t)$) that is extracted from the users. As a result, the additional cost of providing additional utility in the last period of growth does not outweigh the additional benefit from the extra revenue the entrepreneur generates in the exploitation

phase. Thus, the immediate effect on the entrepreneur's revenue of an increase in the size of the locked-in effect is negative.

As a secondary effect, an increase in the size of the locked-in effect u , decreases the maximum size of the network $\bar{\mu}$, as was shown in Lemma 2. Note that the change in the maximum size of the network size is only relevant for the last period of growth and the following period of exploitation, but not for the first periods of network growth. As such, the smaller amount of users that the entrepreneur has to provide utility level δu to in the last period of growth is offset by an equally smaller amount of users that the entrepreneur can exploit by providing utility level $-(1 - \delta)u$ in the following periods. Further, the entrepreneur's revenue is increasing in the size of the network, such that a decrease in the network size decreases the entrepreneur's revenue. As both the immediate and secondary effects on the entrepreneur's revenue from an increase in the size of the locked-in effect are negative, the total effect is negative. Thus, the entrepreneur's revenue with centralized governance is decreasing in u .

1.2.4 Welfare

Finally, I want to address the welfare implications of the governance decision. In particular: When does decentralization improve welfare? It turns out, that this question can be answered with the analysis that has been conducted so far. First, note that users in the centralized implementation of the network are always indifferent between joining the network and their outside option ex-ante. In contrast, users receive strictly positive utility in the decentralized implementation of the network. Thus, users always prefer decentralized governance. For the entrepreneur, proposition 3 has established that she prefers decentralization if and only if the size of the locked in effect u is larger than the threshold u^{**} . Therefore, the following corollary can be established:

Corollary 2 *Decentralized governance of the network is a Pareto improvement over centralized governance if and only if the size of the locked-in effect u is larger than u^{**}*

As an alternative, one might consider utilitarian welfare. Then, utilitarian welfare is also increased through decentralization if decentralization constitutes a Pareto improvement, i.e. if the size of the locked-in effect u is larger than u^{**} . However, the statement for utilitarian welfare is not an if and only if statement. In general, it is not obvious whether it would improve welfare to force an entrepreneur to decentralize her network when locked-in effects are smaller than u^{**} . Doing so creates two welfare effects with opposing signs: the decrease in welfare through the decrease in revenue for the entrepreneur, and the increase in welfare through the increase in utility for the users. The sign of the aggregate of these two effects will generally depend on the parametrization of the model.

1.3 Discussion

1.3.1 Airdrops: Decentralizing at a Later Time

One thing that commonly occurs in practice when an entrepreneur decentralizes her network are so-called *airdrops*. That is, instead of decentralizing her network at the very beginning, the entrepreneur delays decentralizing her network until a later time. At the time of decentralization, past users are then sent tokens to their wallet (the tokens are "air-dropped") and moving forward, the network is subject to decentralized governance.¹⁵ This practice of airdrops can be rationalized within my model by allowing the entrepreneur to delay decentralizing her network until a later period. The main concern is once again commitment, that is, the entrepreneur is unable to commit that she will decentralize the network in the future. Instead, it has to be sequentially optimal for her to decentralize the network. Intuitively, she cannot delay decentralizing for too long, as otherwise exploiting the users becomes too tempting.

Lemma 4 *Suppose $g(\mu_t) - \mu_t \rightarrow 0$ as $\mu_t \rightarrow \infty$. Then, for a sufficiently large network size μ_t it is sequentially optimal to keep the network centralized.*

Proof. See appendix 1.5.10 ■

Next, I show that the entrepreneur can increase her revenues by delaying decentralization of her network for some time. The intuition is, that at the start, when the amount of users in the network is small, a centralized entrepreneur gains implicit commitment to not exploit the locked-in effect of the users by the prospects of future growth. Using that commitment, she can avoid the costs of decentralization for some time to increase her overall profits.

Proposition 4 *Suppose it is optimal for the entrepreneur to decentralize in $t = 0$. Then it is optimal for the entrepreneur to delay decentralizing the network. Further, the option to decentralize the network at a later time increases the range of locked-in effects for which decentralization is optimal.*

Proof. See appendix 1.5.11 ■

As a secondary result, the proposition shows that giving the entrepreneur more flexibility for when she decentralizes, makes decentralization naturally more appealing. Thus, it is profitable for the entrepreneur to (eventually) decentralize her network for even smaller locked-in effects than when she was restricted to decentralizing the network at the very start.

¹⁵For example, Uniswap was founded in November 2018 and decentralized its governance after an airdrop in September 2020.

1.3.2 Equilibrium Multiplicity

Section 2 has discussed the implications of centralized governance for an equilibrium in which the entrepreneur grows the network up to a particular size and then stops growing the network to exploit the locked-in effect of its users. However, there exist other sub-game perfect equilibria.¹⁶ Here, I argue that I perceive them as less convincing, due to the high degree of coordination necessary among the users. In particular, when δ is sufficiently large, there exists the following folk-theorem type of equilibrium:

Users' strategy: Existing users leave the network and newly arriving users do not join the network if the level of utility implied by any revenue sharing α_t and monetization π_t in the history of the game at any time t is strictly lower than the level $\hat{u}_t = V(g(\mu_{t-1})) - (1 - \delta)u$, for a network of size $\mu_t = g(\mu_{t-1})$.

Entrepreneur's strategy: In every period t , set revenue sharing α_t and monetization π_t such that the level of utility for the users is equal to \hat{u}_t for a network of size $\mu_t = g(\mu_{t-1})$. If the entrepreneur is being punished by the users, set utility equal to $-(1 - \delta)u$ conditional on 0 (measure) users being in the network.

A proof that these strategies constitute a sub-game perfect Nash equilibrium can be found in appendix 1.5.9. Now, while this type of equilibrium exists, it is particularly demanding in terms of coordination between the users. To illustrate this point, I will show its instability regarding small uncertainties. Suppose that the entrepreneur deviates and instead offers utility level $\hat{u}_t - \epsilon$ for some arbitrarily small ϵ . Since the utility level of the deviation is arbitrarily close to \hat{u}_t , suppose that user i is not entirely certain whether all other users will follow the equilibrium strategy and punish the entrepreneur by leaving the network/not joining the network. User i assigns probability p to the event that all other users unexpectedly stay in the network, for example because the trigger strategy they follow is slightly more lenient than expected. With probability $1 - p$ all other users leave the network as prescribed by the equilibrium. An equilibrium is considered unstable, if, for a degree of uncertainty of punishment p , there is a small deviation ϵ in the utility offered by the entrepreneur such that any user i is better off staying in the network and not punishing the entrepreneur.

Proposition 5 *The alternative equilibrium discussed in this section is unstable for any degree of uncertainty $p > 0$. In contrast, the equilibrium of the main body of the paper, i.e., in proposition 1, is stable for all degrees of uncertainty.*

Proof. See appendix 1.5.9 ■

¹⁶A multitude of equilibria is common in dynamic games and has been established through various folk theorems (e.g., Abreu (1983), Abreu et al. (1986), Fudenberg and Maskin (1990)).

Intuitively speaking, the folk-theorem style equilibrium has the feature that a particular user i will want to follow through with punishing the entrepreneur for deviating *only if* all other users also follow through. He wants to avoid punishing the entrepreneur, if the other users do not follow suit. Thus, this kind of equilibrium requires an incredibly large degree of coordination. In contrast, the equilibrium presented in the main paper has the feature that a particular user i will want to leave the network (punish the entrepreneur) *regardless* of whether the other users also leave. Thus, no degree of coordination is necessary.

1.4 Conclusion

Before concluding, I want to briefly discuss some further points of interest. First, the reader may wonder if this model implies that an established network such as Google or Facebook should decentralize their business through the blockchain. Such a conclusion cannot be drawn from this model, as these networks have already established a large amount of users (e.g. Facebook already has around 3 billion users¹⁷). As such, the value of extracting additional revenues from existing users that are already locked-in may outweigh the value of commitment that is offered by a decentralized implementation. However, the model provides insights on the optimal governance of newly founded competitors.

Second, it may be plausible that locked-in effects become larger when there are more users. When the network size is small, growth has been shown to be a substitute for commitment in section 1.2.1. Smaller locked-in effects would leave this result unchanged. Further, when the network size, and thus the locked-in effect, would be large, the entrepreneur will find it even more beneficial to stop growing the network and exploit the existing users. Therefore, such an extension will leave the model qualitatively unchanged. Last, consider the possibility that the entrepreneur may treat newly arriving and already existing users differently. For example, she could try to treat newly arriving users or early adopters favorably. However, if this also implies that she can treat existing users less favorably, this change would exacerbate the commitment problem of the entrepreneur when choosing centralized governance even further. That is, it would be sequentially optimal to exploit the locked-in effect of all users as soon as possible. Therefore, commitment should become even more valuable for the entrepreneur.

To summarize, this paper provides an answer to a question that is frequently raised when it comes to the topic of blockchain and cryptocurrencies: *Why should anybody use it?* As the main result, I showed that (i) an entrepreneur prefers to decentralize her net-

¹⁷Meta Earnings Presentation Q2, 2022, p.14

work and (ii) decentralization is a Pareto improvement, if and only if the locked-in effect is sufficiently large. To broaden our understanding of further implications of decentralization, I believe that further research is needed, especially regarding the economics of decentralized governance.

1.5 Appendix

1.5.1 Explanation of Blockchain, Smart Contracts and the Creation of Commitment

This section provides a brief overview over blockchains, smart contracts and some examples of projects that leverage this technology. It is intended to provide sufficient background information for this paper for readers that are not familiar with the topic. However, a thorough treatment of the topic itself is outside the scope of this paper. For a basic introduction to the topic, see for example Lewis (2021). For some further information and more current research, see for example the contributions on <https://www.cberforum.org/>.

Blockchain

A blockchain is a ledger that allows for the storage of information. In this paper, the focus lies on decentralized blockchains, i.e., those that are permissionless, and public. They are updated and maintained decentrally by their users through a consensus mechanism. The two most common consensus mechanisms are Proof of Work and Proof of Stake.¹⁸ For a more detailed introduction to Blockchain, and its consensus mechanisms, see for example Saleh (2021). I focus on the implications of blockchains for economic interactions. As they are permissionless, there is no central authority that can censor access to the blockchain. As such, an entrepreneur that leverages a decentralized blockchain finds herself unable to interfere with the users' ability to use the blockchain. Further, it is tamper-proof, i.e. the entrepreneur and any single user are unable to change records on the blockchain. As the blockchain is public, anyone can publicly observe and trust in the current consensus of information on the blockchain.¹⁹ Bitcoin is probably the most well-known blockchain to date. It was created in 2008 by Satoshi Nakamoto.²⁰ The Bitcoin blockchain securely stores account balances and facilitates transactions between its users.

¹⁸In practice, blockchains are updated by a subset of their users. In Proof of Work blockchains, this subset is commonly referred to as miners. In Proof of Stake blockchains, they are commonly referred to as validators.

¹⁹There are a variety of explorers that allow for easier reading of blockchains. For example, <https://etherscan.io/> covers the Ethereum blockchain.

²⁰Satoshi Nakamoto is a pseudonym. The real name of the bitcoin founder is unknown. Furthermore, it is unknown if Satoshi Nakamoto is a single person or a group of people.

Smart Contracts

From a technological standpoint, Bitcoin is not as advanced as many newer blockchains. Most notably, it is not smart contract compatible.²¹ Essentially, a smart contract is a piece of code that can be executed on the blockchain. Smart contracts have first been formalized by Szabo (1997). The first smart contract compatible blockchain, Ethereum, was conceived in a white paper by Vitalik Buterin in 2014.²² Smart contract compatible blockchains offer vast possibilities for interactions between economic agents in a trustless environment. For example, they can be programmed to facilitate the exchange of cryptocurrencies between two economic agents, without the need for trust in each other or a central party as an intermediary. To date, the top 10 cryptocurrencies by market capitalization consist of Bitcoin, three stablecoins, and six smart contract compatible blockchains.²³ This highlights the growing importance of smart contract compatible blockchains.

Creating commitment through smart contracts: the example of Uniswap

One of the simplest examples of networks that rely on smart contracts to govern the economic interactions between its users is decentralized exchanges. The largest decentralized exchange to date is Uniswap²⁴. It was founded in November 2018 by Hayden Adams and deployed on the Ethereum blockchain. Uniswap allows its users to exchange different cryptocurrencies in a trustless environment using smart contracts as intermediaries. As of September 2022, it has facilitated the exchange of roughly \$1.1 trillion worth of cryptocurrencies in 110 million trades. As the exchange is facilitated by smart contracts, which are immutable once deployed to the blockchain, the terms of the exchange remain unchanged at a 0.3% fee, regardless of how popular it has become.²⁵ It is entirely impossible for Adams to change the terms of the smart contracts governing Uniswap to extract additional rents from its sizeable user base. Changes to the Uniswap protocol are facilitated through a decentralized governance mechanism that uses UNI governance tokens.²⁶ Such an arrangement is also referred to as a *Decentralized Autonomous Organization (DAO)*. Changes to the protocol are then voted on in a majority vote where 1 token equals 1 vote.

²¹As pointed out in the Ethereum white paper, technically Bitcoin can perform some computations, but it is severely limited. For example, it is not *Turing complete*.

²²<https://ethereum.org/en/whitepaper/ethereum>

²³At the date of writing the top 10 cryptocurrencies are: Bitcoin, 3 stablecoins (USDT, USDC, BUSD) and 6 smart contract compatible blockchains (Ethereum, Binance Smart Chain, Ripple, Cardano, Solana and Dogecoin. A stablecoin is a cryptocurrency pegged to a fiat currency, most commonly the US Dollar.

²⁴<https://uniswap.org/>

²⁵There are other Uniswap smart contracts available with fees of 0.01%, 0.05% and 1% respectively.

²⁶The UNI governance token is a digital asset. A digital asset is referred to as a cryptocurrency if it has its own underlying blockchain. If it utilizes another blockchain, it is referred to as a token. UNI exists under the ERC-20 token standard on the Ethereum blockchain

Creating commitment when smart contracts are not sufficient: the example of Presearch

For some networks, it is not feasible to contain the entire interaction between agents within a smart contract. Consider the example of Presearch, a decentralized search engine. When a user searches on a search engine, a simplified workflow is as follows: 1) The user issues a search request and sends it to the search engine, 2) the search engine computes the search results and sends them back to the user. If one were to try to contain this interaction in a smart contract, there would be at least two serious challenges. First, the block creation times on current blockchains range from minutes (Bitcoin) to seconds (Ethereum) to several hundred milliseconds (Solana). As such, the execution of a search through a smart contract would simply be too slow to be practical. Second, interaction with a smart contract requires the user to pay for gas fees²⁷. With Ethereum, these gas fees are typically in the range of several dollars.²⁸ As such, they are too high to facilitate millions to billions of searches a day.²⁹ Therefore, for many networks, at least some interactions have to happen off-chain.

To see how this works in practice, consider an entrepreneur who wishes to create a search engine. In a centralized implementation, she develops the code and sets up a data center with the computing infrastructure to handle the users' search requests. To monetize her search engine, she allows advertisers to place advertisements within the search results. The entrepreneur starts off with minimal advertisement to attract new users. As users are locked-in to her search engine, she increases the number of advertisements she displays with the search results. Suppose the users anticipate this behavior by the entrepreneur and that it is necessary for the entrepreneur to be able to commit. How can she create commitment through decentralization and the blockchain?

Instead of operating the search engine through her own infrastructure, she decides to distribute the code of the search engine freely and asks her users to set up the infrastructure (a so-called node) for the search engine. Now, suppose the entrepreneur tries to update the software to increase the advertisement on the search engine as the users have become locked in. Since the users are effectively operating the search engine, they can simply refuse to install the software update that the entrepreneur has put forward. Thus, the entrepreneur is unable to abuse the locked-in effect of the users. So far, this does not necessarily require the use of the blockchain. However, to compensate the users for the costs of operating the infrastructure, the entrepreneur promises to share part of

²⁷Gas fees are transaction fees that have to be paid to interact with a smart contract on a blockchain. They are necessary to ensure that computations finish within a finite amount of time and keep malicious actors from impeding the operation of the blockchain through endless smart contract calculations.

²⁸Current Ethereum gas fees can be found using <https://etherscan.io/gastracker>

²⁹For example, Google handles around 5-6 billion search requests a day.

the advertisement revenue with them. In this interaction, an opportunity for blockchain technology to mitigate economic frictions arises.

Suppose the entrepreneur has promised the users 50% of advertisement revenues. Further, suppose there are two potential advertisers that are willing to pay \$100 to advertise on the search engine, but it is only possible to display advertisements from one of the advertisers. The willingness to pay is known to the entrepreneur, but not the users. The payment of the advertisers to the entrepreneur is not publicly observable. If everybody behaves honestly, competition will drive the advertisers to pay \$100 for the advertisement, and the entrepreneur and the users will receive \$50 each. Now suppose that the entrepreneur and one of the advertisers decide to collude: The entrepreneur proposes that she will tell the users that the advertiser was only willing to pay \$50 for the advertisement. The other \$50 will be split 30-20 between the entrepreneur and the advertiser. Such collusion between the entrepreneur and the advertiser is profitable for both, since now, the entrepreneur pockets \$55 and the advertiser gets to advertise on the search engine for \$80 instead of \$100. If the users anticipate such collusion, it may be optimal for them to refrain from operating a node in the first place.

This situation can be remedied through the use of the blockchain: when setting up her search engine, the entrepreneur employs a smart contract on the blockchain. It is structured such that advertisers pay the smart contract for the advertisement. The software of the search engine is programmed, such that it displays the advertisement for the highest paying advertiser in the smart contract. Revenues are distributed 50/50 between the entrepreneur and the users using the smart contract. Now collusion between the entrepreneur and one of the advertisers is no longer possible: Suppose the entrepreneur and one of the advertisers agree to pay \$50 for advertising into the smart contract and again split the other \$50 between each other. Now the second advertiser can simply deposit \$51 into the smart contract to have their advertisement displayed, breaking the possibility of collusion between the other advertiser and the entrepreneur. In this example, the decentralized network run by the users serves as a commitment device for the entrepreneur to not abuse their locked-in effect through increased advertisement. The blockchain serves as a commitment device for the entrepreneur to honor her revenue-sharing agreement with the users.

1.5.2 Myopic Revenue Maximization

Lemma 5 *Consider the entrepreneur's problem to maximize revenue myopically in a single period t while ensuring utility \hat{u} for users when the network size is μ_t .*

1. If $\frac{\phi(\mu_t)^2}{4k\mu_t^2} + V(\mu_t) < \hat{u}$ the entrepreneur is unable to ensure utility \hat{u} for the users.

2. If $\frac{\phi(\mu_t)^2}{4k\mu_t^2} + V(\mu_t) \geq \hat{u}$ and

(a) $\left(\frac{\phi(\mu_t)}{2k\mu_t}\right)^2 \geq \frac{V(\mu_t) - \hat{u}}{k}$, the optimal π_t, α_t are given by

$$\pi_t = \frac{\phi(\mu_t)}{2k\mu_t} \quad (1.5.1)$$

$$\alpha_t = \frac{1}{2} + \frac{2k\mu_t^2(V(\mu_t) - \hat{u})}{\phi(\mu_t)^2} \quad (1.5.2)$$

The entrepreneur's revenue is equal to

$$\mu_t V(\mu_t) + \frac{\phi(\mu_t)^2}{4k\mu_t} - \mu_t \hat{u} \quad (1.5.3)$$

(b) $\left(\frac{\phi(\mu_t)}{2k\mu_t}\right)^2 < \frac{V(\mu_t) - \hat{u}}{k}$, the optimal π_t, α_t are given by

$$\pi_t = \sqrt{\frac{V(\mu_t) - \hat{u}}{k}} \quad (1.5.4)$$

$$\alpha_t = 1 \quad (1.5.5)$$

The entrepreneur's revenue is equal to

$$\sqrt{\frac{V(\mu_t) - \hat{u}}{k}} \phi(\mu_t) \quad (1.5.6)$$

Proof. The lemma follows from the following maximization problem:

$$\max_{\alpha_t \pi_t} \alpha_t \pi_t \phi(\mu_t) \quad (1.5.7)$$

$$\text{s.t. } V(\mu_t) - k\pi_t^2 + \frac{1 - \alpha_t}{\mu_t} \pi_t \phi(\mu_t) = \hat{u} \quad (1.5.8)$$

$$1 \geq \alpha_t \geq 0 \quad (1.5.9)$$

The problem can be solved through a standard KKT approach. The FOCs associated

with the resulting Lagrangian with the complementary slackness conditions then reads

$$\frac{\partial}{\partial \alpha_t} = \pi_t \phi(\mu_t) + \lambda_1 \left(\frac{-\pi_t}{\mu_t} \phi(\mu_t) \right) - \lambda_2 + \lambda_3 = 0 \quad (1.5.10)$$

$$\frac{\partial}{\partial \pi_t} = \alpha_t \phi_t(\mu_t) + \lambda_1 \left(-2k\pi_t + \frac{1 - \alpha_t}{\mu_t} \phi(\mu_t) \right) = 0 \quad (1.5.11)$$

$$\frac{\partial}{\partial \lambda_1} = V(\mu_t) - k\pi_t^2 + \frac{1 - \alpha_t}{\mu_t} \pi_t \phi(\mu_t) - \hat{u} = 0 \quad (1.5.12)$$

$$\frac{\partial}{\partial \lambda_2} \lambda_2 = (1 - \alpha_t) \lambda_2 = 0 \quad (1.5.13)$$

$$\frac{\partial}{\partial \lambda_3} \lambda_3 = \alpha_t \lambda_3 = 0 \quad (1.5.14)$$

First, focus on the case where $\alpha_t \in (0, 1)$, such that $\lambda_2, \lambda_3 = 0$. Then straightforward calculations yield that

$$\pi_t = \frac{\phi(\mu_t)}{2k\mu_t} \quad (1.5.15)$$

$$\alpha_t = \frac{1}{2} + \frac{2k\mu_t^2(V(\mu_t) - \hat{u})}{\phi(\mu_t)^2} \quad (1.5.16)$$

And the entrepreneur's revenue equals

$$\left(V(\mu_t) + \frac{\phi(\mu_t)}{4k\mu_t^2} - \hat{u} \right) \mu_t \quad (1.5.17)$$

Note that $\alpha_t \in (0, 1)$ requires that

$$\alpha_t > 0 \quad (1.5.18)$$

$$\iff \frac{\phi(\mu_t)^2}{4k\mu_t^2} + V(\mu_t) \geq \hat{u} \quad (1.5.19)$$

and

$$1 > \alpha_t \quad (1.5.20)$$

$$\iff \left(\frac{\phi(\mu)}{2k\mu_t} \right)^2 > \frac{V(\mu_t) - \hat{u}}{k} \quad (1.5.21)$$

Next, consider the possible solution with $\alpha_t = 1$. Then it follows that

$$\pi_t = \sqrt{\frac{V(\mu_t) - \hat{u}}{k}} \quad (1.5.22)$$

The entrepreneur's revenue then equals

$$\sqrt{\frac{V(\mu_t) - \hat{u}}{k}} \phi(\mu_t) \quad (1.5.23)$$

Last, consider the possible solution where $\alpha_t = 0$. Notice that in this case, the entrepreneur's revenue is equal to 0, regardless of the choice of π_t . The choice of π_t that maximizes the users' utility is $\pi_t = \frac{\phi(\mu_t)}{2k\mu_t}$. Then it is not possible to ensure utility \hat{u} for the user if

$$V(\mu_t) - k \left(\frac{\phi(\mu_t)}{2k\mu_t} \right)^2 + \frac{\phi(\mu_t) \phi(\mu_t)}{2k\mu_t \mu_t} < \hat{u} \quad (1.5.24)$$

$$\iff V(\mu_t) + \frac{\phi(\mu_t)^2}{4k\mu_t^2} < \hat{u} \quad (1.5.25)$$

■

1.5.3 Sufficient conditions for $\bar{\mu}$ to be well-defined

In this section I first provide sufficient conditions for the existence and uniqueness of $\bar{\mu}$ and then discuss how these conditions can be weakened further. Consider the following conditions:

For existence:

1. As $\mu_t \rightarrow \infty$ it holds that $g(\mu_t) - \mu_t \rightarrow 0$
2. $\psi(g(\mu_t), \hat{u}) - \psi(\mu_t, \hat{u})$ is decreasing in μ_t for all \hat{u}

For uniqueness:

1. $\sqrt{2kV'(\mu_t)\mu_t} < \frac{\phi(\mu_t)}{\mu_t}$ for all $\mu_t > 0$

First, I provide an intuitive description of the conditions. They represent the idea that user growth will slow down over time and there are decreasing returns to the entrepreneur's revenue when growing the network. As the size of the network increases, fewer new users will arrive. This condition should be satisfied in many applications, as the potential amount of users of a network is limited. Further, the conditions impose a regularity on the difference between the revenue that the entrepreneur generates. As the network grows, the gap between the revenue created from a network that has grown one more time and a network that has not, shrinks.

Mathematically, the condition requires that the revenue function ψ , which depends on the functions V and ϕ , is not too convex in the network size μ_t , in relation to the rate at which the network growth slows down over time. To illustrate the point, consider an

example with $V(\mu_t)$ constant and $\phi(\mu_t) = \mu_t$. Note that in this specification ψ is a linear function in μ_t , and as such at the extreme end of not too convex functions that can be considered.

Proposition 6 *The conditions presented above are sufficient to guarantee the existence and uniqueness of $\bar{\mu}$.*

Proof of proposition 6

Recall the definition of $\bar{\mu}$ as the value that solves the equation

$$\frac{1}{1-\delta}\psi(\bar{\mu}, -(1-\delta)u) = \psi(g(\bar{\mu}), \delta u) + \frac{\delta}{1-\delta}\psi(g(\bar{\mu}), -(1-\delta)u) \quad (1.5.26)$$

Note that at $\mu = 0$ it holds that LHS of equation $<$ RHS of the equation. Evaluating at $\mu \rightarrow \infty$ implies LHS of equation $>$ RHS of the equation. Given the continuity of all functions involved, an application of the intermediate value theorem implies existence. To show the unique cutoff, consider the first derivative of the difference of the RHS and the LHS with respect to μ :

$$g'(\mu)\psi_\mu(g(\mu), \delta u) + g'(\mu)\frac{\delta}{1-\delta}\psi_\mu(g(\mu), -(1-\delta)u) - \frac{1}{1-\delta}\psi_\mu(\mu, -(1-\delta)u) \quad (1.5.27)$$

$$= g'(\mu)\psi_\mu(g(\mu), \delta u) - \psi_\mu(\mu, -(1-\delta)u) + \frac{\delta}{1-\delta}(g'(\mu)\psi_\mu(g(\mu), -(1-\delta)u) - \psi_\mu(\mu, -(1-\delta)u)) \quad (1.5.28)$$

What is to be shown is that this first derivative is negative. To this end, I show the intermediate result that under the assumption that $\sqrt{2kV'(\mu)\mu} < \frac{\phi(\mu)}{\mu}$ for all $\mu > 0$ it holds that $\frac{\partial\psi^2}{\partial\mu\partial\hat{u}} < 0$ for all $\mu > 0$.

Lemma 6 $\sqrt{2kV'(\mu)\mu} < \frac{\phi(\mu)}{\mu}$ for all $\mu > 0$ implies $\frac{\partial\psi^2}{\partial\mu\partial\hat{u}} < 0$ for all $\mu > 0$.

Proof. Note that

$$\frac{\partial\psi^2}{\partial\mu\partial\hat{u}} = \begin{cases} -1 & \text{if } \left(\frac{\phi(\mu)}{2k\mu}\right)^2 \geq \frac{V(\mu)-\hat{u}}{k} \\ -\frac{\phi'(\mu)}{2\sqrt{k}}(V(\mu)-\hat{u})^{-0.5} + \frac{V'(\mu)}{4\sqrt{k}}(V(\mu)-\hat{u})^{-1.5}\phi(\mu) & \text{if } \left(\frac{\phi(\mu)}{2k\mu}\right)^2 < \frac{V(\mu)-\hat{u}}{k} \end{cases} \quad (1.5.29)$$

Therefore I focus on showing that the second case is negative:

$$-\frac{\phi'(\mu)}{2\sqrt{k}}(V(\mu)-\hat{u})^{-0.5} + \frac{V'(\mu)}{4\sqrt{k}}(V(\mu)-\hat{u})^{-1.5}\phi(\mu) < 0 \quad (1.5.30)$$

$$\iff -2\phi'(\mu)(V(\mu)-\hat{u}) + V'(\mu)\phi(\mu) < 0 \quad (1.5.31)$$

Note that to be in this second case, \hat{u} is bounded above such that $\hat{u} < -\left(\frac{\phi(\mu)}{2k\mu}\right)^2 k + V(\mu)$. Therefore, it holds that

$$-2\phi'(\mu)(V(\mu) - \hat{u}) + V'(\mu)\phi(\mu) < -2\phi'(\mu)\left(\frac{\phi(\mu)}{2k\mu}\right)^2 k + V'(\mu)\phi(\mu) \quad (1.5.32)$$

This is smaller than 0 if

$$-2\phi'(\mu)\left(\frac{\phi(\mu)}{2k\mu}\right)^2 k + V'(\mu)\phi(\mu) < 0 \quad (1.5.33)$$

$$\iff 2k\mu^3 V'(\mu) \frac{\phi(\mu)}{\phi'(\mu)\mu} < \phi(\mu)^2 \quad (1.5.34)$$

Note that the assumption that $\frac{\phi(\mu_t)}{\mu_t}$ is non-decreasing guarantees that $\frac{\phi(\mu)}{\phi'(\mu)\mu} \leq 1$. This implies that the inequality below is a sufficient condition for 1.5.34

$$\sqrt{2kV'(\mu)\mu} < \frac{\phi(\mu)}{\mu} \quad (1.5.35)$$

Which is the uniqueness part of the conditions. ■

Now, I revisit the initial derivative

$$g'(\mu)\psi_\mu(g(\mu), \delta u) - \psi_\mu(\mu, -(1-\delta)u) + \frac{\delta}{1-\delta}(g'(\mu)\psi_\mu(g(\mu), -(1-\delta)u) - \psi_\mu(\mu, -(1-\delta)u)) \quad (1.5.36)$$

Using the lemma derived above, note that $\psi_\mu(\mu, \delta u) < \psi_\mu(\mu, -(1-\delta)u)$. Thus, it holds hat

$$g'(\mu)\psi_\mu(g(\mu), \delta u) - \psi_\mu(\mu, -(1-\delta)u) + \frac{\delta}{1-\delta}(g'(\mu)\psi_\mu(g(\mu), -(1-\delta)u) - \psi_\mu(\mu, -(1-\delta)u)) \quad (1.5.37)$$

$$< g'(\mu)\psi_\mu(g(\mu), \delta u) - \psi_\mu(\mu, \delta u) + \frac{\delta}{1-\delta}(g'(\mu)\psi_\mu(g(\mu), -(1-\delta)u) - \psi_\mu(\mu, -(1-\delta)u)) \quad (1.5.38)$$

Further, the assumption that $\psi(g(\mu_t), \hat{u}) - \psi(\mu_t, \hat{u})$ is decreasing in μ_t for all \hat{u} implies that

$$g'(\mu)\psi_\mu(g(\mu), \hat{u}) - \psi_\mu(\mu, \hat{u}) \leq 0 \quad (1.5.39)$$

Using this implies that expression 1.5.38 is smaller than 0 which finishes the proof.

An example with a general growth function and linear revenues:

First off, I show that the specification of $V(\mu)$ constant and $\phi(\mu) = \mu$ with $g(\mu) = \mu + \gamma(\mu)$ and γ being strictly decreasing, strictly positive and approaching 0 as $\mu \rightarrow \infty$ satisfy the sufficient conditions above. Clearly, as $\mu \rightarrow \infty$ it holds that $g(\mu_t) - \mu_t \rightarrow 0$ as $\gamma(\mu) \rightarrow 0$ as $\mu \rightarrow \infty$. Next, consider the difference $\psi(g(\mu_t), \hat{u}) - \psi(\mu_t, \hat{u})$. Plugging in V and ϕ yields that $\psi(\mu_t, \hat{u})$ is a linear function of μ_t . Now for the assumption to hold, consider the first derivative of the difference $\psi(g(\mu_t), \hat{u}) - \psi(\mu_t, \hat{u})$:

$$\frac{\partial}{\partial \mu_t} (\psi(g(\mu_t), \hat{u}) - \psi(\mu_t, \hat{u})) = g'(\mu_t) \psi_{\mu_t}(g(\mu_t), \hat{u}) - \psi_{\mu_t}(\mu_t, \hat{u}) \quad (1.5.40)$$

$$= g'(\mu_t) \psi_{\mu_t}(\mu_t, \hat{u}) - \psi_{\mu_t}(\mu_t, \hat{u}) \quad (1.5.41)$$

$$= \gamma'(\mu_t) \psi_{\mu_t}(\mu_t, \hat{u}) < 0 \quad (1.5.42)$$

The condition for uniqueness can be easily confirmed.

An example with a general revenue function and growth that slows abruptly:

For another example, consider the opposite end of the spectrum. That is, consider a growth function $g(\mu)$ such that

$$g(\mu) = \begin{cases} g(0) > 0 & \text{if } \mu = 0 \\ \mu & \text{if } \mu > 0 \end{cases} \quad (1.5.43)$$

and arbitrary functions $V(\mu)$ and $\phi(\mu)$. Then clearly we have

$$\psi(g(\mu), \hat{u}) - \psi(\mu, \hat{u}) > 0 \quad (1.5.44)$$

if $\mu = 0$ and the difference equals 0 otherwise. Intuitively speaking, this growth function allows the network to grow for exactly 1 period at the start, and then in future periods no new users arrive. Restricting the growth function in this way allows for maximum freedom regarding the functions V and ϕ .³⁰

To recap, the sufficient conditions rely on a balance between the convexity of the revenue function ψ in relation to the growth function g . For the minimum degree of convexity of ψ , i.e., when ψ is linear when V is constant and $\phi(\mu) = \mu$ it is possible to allow very general growth functions g . On the other end, it is possible to allow very general functions V and ϕ , implying very general shapes on the revenue function ψ , if growth slows down extremely fast, that is, decreases to 0 within 1 period. In general, appropriate

³⁰Note that this definition of g includes a discontinuity. To use such a g in the model, one would have to extend g to a continuous function or use a slightly more general definition of $\bar{\mu}$, both of which can be accommodated fairly easily.

functions for V , ϕ and g can be found by keeping in mind the trade-off between relatively more convex revenue functions ψ (as calculated by V and ϕ) for growth functions g that slow down relatively faster and vice-versa.

More general sufficient conditions:

What is important for the proofs in the paper is that $\bar{\mu}$ exists and is unique. For this, I have presented sufficient conditions above. However, they are not necessary. Alternatively, it is possible to assume that

$$\psi(g(\mu), \delta u) + \frac{\delta}{1-\delta} \psi(g(\mu), -(1-\delta)u) - \frac{1}{1-\delta} \psi(\mu, -(1-\delta)u) \quad (1.5.45)$$

is

1. Increasing up to some value $\tilde{\mu}$
2. Strictly decreasing for any $\mu > \tilde{\mu}$

This case carries the intuition that the network effects through the entry of additional users outweigh a slowdown in growth up to $\tilde{\mu}$ users. Afterwards, the relationship reverses. Note that mathematically this assumption also guarantees the existence of a unique $\bar{\mu}$ and that it is more general in the sense that it contains the sufficient conditions from above for the case where $\tilde{\mu} = 0$. However, it is considerably more challenging to calculate examples that satisfy this assumption.

1.5.4 Proof of Proposition 1

To check for profitable deviations by the entrepreneur or the users, I employ the one-shot deviation principle (see for example Theorem 4.2 in Fudenberg and Tirole (1991)). Note that the one-shot deviation principle applies, as the game is obviously continuous at infinity.³¹ Therefore, it is sufficient to check that there is no single period profitable deviation.

Deviations by the entrepreneur:

Consider a history of the game up to some period t that results in a mass of users μ_{t-1} at the start of the period. Then there are two cases:

³¹c.f. Definition 4.1 and explanation in Fudenberg and Tirole (1991): A game is continuous at infinity if for each player i the utility function u_i satisfies $\sup_{h, \tilde{h} \text{ s.t. } h^t = \tilde{h}^t} |u_i(h) - u_i(\tilde{h})| \rightarrow 0$ as $t \rightarrow \infty$. It is satisfied if the overall payoffs are a discounted sum of per-period payoffs and the per period payoffs are uniformly bounded.

Case 1 ($\mu_{t-1} \leq \bar{\mu}$): First, note that deviations that increase the utility of the users are not profitable, since the equilibrium path remains unchanged and the entrepreneur's revenue is decreasing in the utility level she provides to the users. Now, consider a deviation that decreases the utility level the entrepreneur provides for the users. Given the users' strategies, a large decrease in the utility level below $-(1 - \delta)u$ will cause all users to leave the network and not be profitable. A small decrease will cause existing users to remain in the network and newly arriving users to not join the network. Therefore, the most profitable deviation would be to a utility level of $-(1 - \delta)u$. The entrepreneur's revenue for this deviation is $\psi(\mu_{t-1}, -(1 - \delta)u)$ plus the discounted revenue of the continuation of the initial strategy starting in the next period. If the entrepreneur had not deviated, she would receive the value of the continuation of the initial strategy starting this period. Note that this value depends on how many more periods the entrepreneur will grow the network according to the initial strategy. I show that the deviation is not profitable by induction on the number of periods of future growth. First, consider the case with 1 period of future growth. Then the deviation is not profitable if

$$\psi(\mu_{t-1}, -(1 - \delta)u) + \delta \left(\psi(g(\mu_{t-1}), \delta u) + \frac{\delta}{1 - \delta} \psi(g(\mu_{t-1}), -(1 - \delta)u) \right) \quad (1.5.46)$$

$$\leq \psi(g(\mu_{t-1}), \delta u) + \frac{\delta}{1 - \delta} \psi(g(\mu_{t-1}), -(1 - \delta)u) \quad (1.5.47)$$

$$\iff \frac{1}{1 - \delta} \psi(\mu_{t-1}, -(1 - \delta)u) \leq \psi(g(\mu_{t-1}), \delta u) + \frac{\delta}{1 - \delta} \psi(g(\mu_{t-1}), -(1 - \delta)u) \quad (1.5.48)$$

Which holds true since $\mu_{t-1} \leq \bar{\mu}$. Now suppose that it is not profitable to deviate when there are T periods of future growth. Next, I show that it is not profitable to deviate with $T + 1$ periods of future growth. A deviation with $T + 1$ periods of future growth is not profitable if

$$\psi(\mu_{t-1}, -(1 - \delta)u) + \delta \left(\sum_{s=0}^{T-1} \delta^s \psi(g^{(s)}(\mu_{t-1}), 0) + \delta^T \psi(g^{(T)}(\mu_{t-1}), \delta u) + \frac{\delta^{T+1}}{1 - \delta} \psi(g^{(T)}(\mu_{t-1}), -(1 - \delta)u) \right) \quad (1.5.49)$$

$$\leq \sum_{s=0}^{T-1} \delta^s \psi(g^{(s)}(\mu_{t-1}), 0) + \delta^T \psi(g^{(T)}(\mu_{t-1}), \delta u) + \frac{\delta^{T+1}}{1 - \delta} \psi(g^{(T)}(\mu_{t-1}), -(1 - \delta)u) \quad (1.5.50)$$

$$\iff \frac{1}{1 - \delta} \psi(\mu_{t-1}, -(1 - \delta)u) \quad (1.5.51)$$

$$\leq \sum_{s=0}^{T-1} \delta^s \psi(g^{(s)}(\mu_{t-1}), 0) + \delta^T \psi(g^{(T)}(\mu_{t-1}), \delta u) + \frac{\delta^{T+1}}{1 - \delta} \psi(g^{(T)}(\mu_{t-1}), -(1 - \delta)u) \quad (1.5.52)$$

Since by induction the assertion holds true for T periods of future growth, it suffices to show that the RHS of the inequality above for T periods of future growth is smaller than the RHS of the inequality above for $T + 1$ periods of future growth, since the LHS is

identical in both cases. Thus, I have to show that

$$\sum_{s=0}^{T-2} \delta^s \psi(g^{(s)}(\mu_{t-1}), 0) + \delta^{T-1} \psi(g^{(T-1)}(\mu_{t-1}), \delta u) + \frac{\delta^T}{1-\delta} \psi(g^{(T-1)}(\mu_{t-1}), -(1-\delta)u) \quad (1.5.53)$$

$$\leq \sum_{s=0}^{T-1} \delta^s \psi(g^{(s)}(\mu_{t-1}), 0) + \delta^T \psi(g^{(T)}(\mu_{t-1}), \delta u) + \frac{\delta^{T+1}}{1-\delta} \psi(g^{(T)}(\mu_{t-1}), -(1-\delta)u) \quad (1.5.54)$$

$$\iff \delta^{T-1} \psi(g^{(T-1)}(\mu_{t-1}), \delta u) + \frac{\delta^T}{1-\delta} \psi(g^{(T-1)}(\mu_{t-1}), -(1-\delta)u) \quad (1.5.55)$$

$$\leq \delta^{T-1} \psi(g^{(T-1)}(\mu_{t-1}), 0) + \delta^T \psi(g^{(T)}(\mu_{t-1}), \delta u) + \frac{\delta^{T+1}}{1-\delta} \psi(g^{(T)}(\mu_{t-1}), -(1-\delta)u) \quad (1.5.56)$$

Now note that $\psi(g^{(T-1)}(\mu_{t-1}), \delta u) < \psi(g^{(T-1)}(\mu_{t-1}), 0)$. Then this implication and some rearranging yields

$$\frac{1}{1-\delta} \psi(g^{(T-1)}(\mu_{t-1}), -(1-\delta)u) \leq \psi(g^{(T)}(\mu_{t-1}), \delta u) + \frac{\delta}{1-\delta} \psi(g^{(T)}(\mu_{t-1}), -(1-\delta)u) \quad (1.5.57)$$

Which holds true since this is precisely the condition that it is optimal to grow $T+1$ times. Therefore, one-shot deviations by the entrepreneur to abuse the locked-in effect of the users are not profitable.

Case 2 ($\mu_{t-1} > \bar{\mu}$): For this case, deviations that decrease the user utility are not profitable, since they will result in all users leaving the network and zero revenues. Now consider deviations that increase the users' utility. First, marginal increases will not change the user behavior on the equilibrium path and are not profitable. Second, the smallest deviation that changes the users' behavior on the equilibrium path is to increase the utility sufficiently to grow the network one more time. However, by definition of $\bar{\mu}$ such deviations are not profitable when $\mu_{t-1} > \bar{\mu}$.

Deviations by a user:

Newly arriving users: First, consider any histories on the equilibrium path. Then, there is no profitable deviation, since users are exactly indifferent between joining and not joining the network. Now, consider deviations off the equilibrium path. For any histories that offer more utility than the equilibrium path, clearly it is still optimal to join the network, such that not joining is not a profitable deviation. In contrast, any histories that

have reduced utility imply that it is optimal to not join the network, such that joining is not a profitable deviation.

Users that are locked-in: First, consider any histories on the equilibrium path. There are two cases. Before the exploitation phase begins, there are no profitable deviations since remaining in the network provides 0 utility, while leaving gives utility $-u < 0$. During the exploitation phase, the users are indifferent between staying and leaving, such that leaving is not a profitable deviation.

Second, consider histories off the equilibrium path. Histories that result in increased user utility obviously do not offer profitable deviations. Now, consider histories such that the user's utility is reduced. Leaving the network provides $-u$ utility, while remaining in the network provides the user a utility level smaller than $-(1-\delta)u$ for the period in which he is alone in the network and utility $-\delta u$ from leaving the network the next period. Total utility is thus smaller than $-(1-\delta)u - \delta u = -u$, such that the deviation is not profitable.

1.5.5 Proof of Lemma 2

First, I show that the implicit function theorem is applicable in this situation. In particular, it has to be shown that the revenue function is differentiable. Clearly, it is piece-wise differentiable. However, it has to be shown that it is also differentiable at the point where the entrepreneur stops revenue sharing, i.e., when

$$\left(\frac{\phi(\mu_t)}{2k\mu_t}\right)^2 = \frac{V(\mu_t) - \hat{u}}{k} \quad (1.5.58)$$

The two pieces of the function are

$$\mu_t V(\mu_t) + \frac{\phi(\mu_t)^2}{4k\mu_t} - \mu_t \hat{u} \quad (1.5.59)$$

and

$$\sqrt{\frac{V(\mu_t) - \hat{u}}{k}} \phi(\mu_t) \quad (1.5.60)$$

Consider differentiability regarding \hat{u} . The derivatives regarding \hat{u} are

$$-\mu_t \quad (1.5.61)$$

and

$$-\frac{1}{2\sqrt{k}} \frac{1}{\sqrt{V(\mu_t) - \hat{u}}} \phi(\mu_t) \quad (1.5.62)$$

It is straightforward to verify algebraically that the two derivatives are equal to each other when $\left(\frac{\phi(\mu_t)}{2k\mu_t}\right)^2 = \frac{V(\mu_t) - \hat{u}}{k}$

Next, I consider the derivatives regarding μ_t . They are

$$V(\mu_t) + \mu_t V'(\mu_t) + \frac{2\phi'(\mu_t)\phi(\mu_t)4k\mu_t - 4k\phi(\mu_t)^2}{(4k\mu_t)^2} - \hat{u} \quad (1.5.63)$$

and

$$\frac{1}{\sqrt{k}} \left(\frac{V'(\mu_t)}{2} \frac{1}{\sqrt{V(\mu_t) - \hat{u}}} \phi(\mu_t) + \sqrt{V(\mu_t) - \hat{u}} \phi'(\mu_t) \right) \quad (1.5.64)$$

Using the identity $\left(\frac{\phi(\mu_t)}{2k\mu_t}\right)^2 = \frac{V(\mu_t) - \hat{u}}{k}$ at the point of interest we can simplify the two derivatives to

$$\left(\frac{\phi(\mu_t)}{2k\mu_t}\right)^2 k + \mu_t V'(\mu_t) + \frac{\phi'(\mu_t)\phi(\mu_t)}{2k\mu_t} - \left(\frac{\phi(\mu_t)}{2k\mu_t}\right)^2 k = \mu_t V'(\mu_t) + \frac{\phi'(\mu_t)\phi(\mu_t)}{2k\mu_t} \quad (1.5.65)$$

and

$$\frac{1}{\sqrt{k}} \left(\frac{V'(\mu_t)}{2\sqrt{k}} \frac{2k\mu_t}{\phi(\mu_t)} \phi(\mu_t) + \frac{\phi(\mu_t)\sqrt{k}}{2k\mu_t} \phi'(\mu_t) \right) = \mu_t V'(\mu_t) + \frac{\phi'(\mu_t)\phi(\mu_t)}{2k\mu_t} \quad (1.5.66)$$

respectively, which are equal to each other. Therefore, the implicit function theorem applies. To shorten notation define

$$F := \psi(g(\mu_t), \delta u) + \frac{\delta}{1-\delta} \psi(g(\mu_t), -(1-\delta)u) - \frac{1}{1-\delta} \psi(\mu_t, -(1-\delta)u) \quad (1.5.67)$$

and by the implicit function theorem it holds that

$$\frac{\partial \bar{\mu}}{\partial u} = - \frac{\frac{\partial F}{\partial u}}{\frac{\partial F}{\partial \mu_t}} \Bigg|_{\bar{\mu}, u} \quad (1.5.68)$$

For the denominator, notice that the derivative is negative by the definition of $\bar{\mu}$.

For the numerator, notice that at $u = 0$ it holds that $F > 0$. Moreover, note that the three parts of F are decreasing and concave, increasing and concave, and decreasing and convex with respect to u respectively. In order for F to be equal to 0 at $(\bar{\mu}, u)$, the derivative of F regarding u has to be negative for at least some values of U . However, note that when the derivative of F turns negative, it will remain negative. This holds, as the middle part of F is increasing and concave, such that its growth slows down. When the derivative turns negative, the third part of F , $-\frac{1}{1-\delta} \psi(\mu_t, -(1-\delta)u)$ alone will keep the derivative negative, as $\mu_t < g(\mu_t)$. Thus, the negative slope is steeper than the positive

slope of $\frac{\delta}{1-\delta}\psi(g(\mu_t), -(1-\delta)u)$. In particular, this implies that the slope of F regarding u at $(\bar{\mu}, u)$ is negative. Therefore, the numerator is negative and the fraction as a whole is negative.

1.5.6 Proof of Proposition 2

The degree of monetization follows from a simple optimization problem. Namely,

$$\max_{\pi_t} V(\mu_t) - k\pi_t^2 + \frac{1-\alpha}{\mu_t}\pi\phi(\mu_t) \quad (1.5.69)$$

The equilibrium is confirmed by an application of the one-shot deviation principle. First, no user has an incentive to deviate in the degree of monetization in weakly dominant strategies. Second, as all users receive strictly positive utility from participation in the network, there is no incentive to deviate into not joining.

Last, the entrepreneur's optimization problem in $t = 0$ equals

$$\max_{\alpha} \sum_{t=1}^{\infty} \left(\alpha \frac{1-\alpha}{2k} \frac{\phi(g^{(t)}(\mu_0))}{g^{(t)}(\mu_0)} \phi(g^{(t)}(\mu_0)) \right) \quad (1.5.70)$$

Where $g^{(t)}$ denotes the t -time chaining of the growth function. From this, it is straightforward to derive $\alpha^* = 0.5$

1.5.7 Proof of Lemma 3

Note that at $u = 0$ the strategy of the entrepreneur is to ensure 0 utility for the users in every period. Further, it holds that there is no value of $\bar{\mu}$ that makes the entrepreneur indifferent between growing the network once more and exploiting the users in the future and exploiting the users right away. Namely, it will always be better to grow the network as $g(\mu) - \mu \geq 0$. Therefore, at $u = 0$ the network will grow every period, as it does with decentralized governance. However, since the choice set regarding monetization and revenue sharing is larger in centralized governance than it is in decentralized governance, her revenues are necessarily higher with centralized governance. Since the entrepreneur's revenues are continuous in u , this result also holds for $u > 0$, but sufficiently close to 0.

1.5.8 Proof of Proposition 3

Corollary 1 established that decentralized governance is preferred over centralized governance if u is sufficiently large, i.e. $u > u^*$. Further, lemma 3 established that centralized governance is preferred if u is sufficiently small. To derive the result of the proposition, note that the entrepreneur's revenue with decentralized governance is independent of

u . Thus, it is sufficient to show that centralized revenue is decreasing in u to prove the proposition. Now, consider the change in the entrepreneur's revenue with centralized governance as u increases. Note that the entrepreneur does not exploit the locked-in effect in the first periods of growth, that is, she sets $\hat{u}_t = 0$ for all periods of growth except the last period. Now, consider the last period of growth and the following periods of exploiting the locked-in effect. Note that the size of the network in all of those periods is the same. Then the first order effect from increasing the size of the locked-in effect is equal to

$$\delta\psi_u(\mu, \delta u) - (1 - \delta)\frac{\delta}{(1 - \delta)}\psi_u(\mu, -(1 - \delta)u) \quad (1.5.71)$$

This is negative if

$$\psi_u(\mu, \delta u) \leq \psi_u(\mu, -(1 - \delta)u) \quad (1.5.72)$$

Now there are three options to compare. They are 1) both sides of the equation are in the linear part of ψ . 2) The LHS is in the linear part and the RHS is in the concave part of ψ . 3) Both sides are in the concave part of ψ . The first case holds trivially. The second case holds as

$$-\mu \leq -\frac{1}{2\sqrt{k}}\frac{1}{\sqrt{V(\mu) + (1 - \delta)u}}\phi(\mu) \quad (1.5.73)$$

$$\iff \left(\frac{\phi(\mu)}{2k\mu}\right)^2 < \frac{V(\mu) + (1 - \delta)u}{k} \quad (1.5.74)$$

Which is a true statement, as it is precisely the condition from lemma 5 that ensured that the RHS is in the concave part of the function.

Last, I show that the inequality holds if both the RHS and the LHS of the equation are in the concave part of ψ .

$$-\frac{1}{2\sqrt{k}}\frac{1}{\sqrt{V(\mu) - \delta u}}\phi(\mu) \leq -\frac{1}{2\sqrt{k}}\frac{1}{\sqrt{V(\mu) + (1 - \delta)u}}\phi(\mu) \quad (1.5.75)$$

$$\iff u \geq 0 \quad (1.5.76)$$

For the second order effect, note that the maximum network size $\bar{\mu}$ is dependent on u . In particular, lemma 2 showed that $\bar{\mu}$ is decreasing in u . Further, the entrepreneur's revenue ψ is increasing in μ , such that the decrease in the maximum size of the network decreases the entrepreneur's revenues. Thus, the total effect of an increase in u on the entrepreneur's revenues is negative.

1.5.9 Proof of equilibrium of section 1.3.2 and proof of proposition 5

Proof of equilibrium of section 1.3.2

First, consider why these strategies constitute a sub-game perfect equilibrium by checking for one shot deviations.

Deviations by the entrepreneur: Given the users strategies, and the fact that the entrepreneur's revenue is decreasing in \hat{u}_t , clearly there are no profitable deviations for the entrepreneur. Increasing \hat{u}_t lowers her revenue without changing the users' behavior on the equilibrium path. Decreasing \hat{u}_t causes all users to leave the network, resulting in 0 revenues for the entrepreneur. When the entrepreneur is being punished and there are no users in the network, the entrepreneur is indifferent between all of his choices, such that there is no incentive to deviate.

Deviations by the users: Fix the strategies of the entrepreneur and the users. Now consider some arbitrary user i . For sub-game perfection, the user cannot have any incentive to (one-shot) deviate from the equilibrium strategy at any history of the game.

First, consider histories of the game such that the entrepreneur has offered at least utility level \hat{u}_t in every period. Suppose user i is already locked into the network. If user i leaves, his utility will be equal to $-u$. If he stays, his utility will be equal to $\sum_{t=0}^{\infty} (\delta^t V(g^{(t)}(\mu_t))) - u$ which is larger than $-u$, such that leaving is not a profitable deviation. Now consider the case where user i is newly arriving to the network. Again, his utility is $\sum_{t=0}^{\infty} (\delta^t V(g^{(t)}(\mu_t))) - u$. This will be larger than 0 for δ large enough, such that there is no incentive to deviate.

Next, consider histories of the game such that the entrepreneur is offering a utility level $\tilde{u}_t < \hat{u}_t$ in some period t . If user i leaves, his utility will be equal to $-u$. If user i stays on the other hand, his utility will be equal to

$$\tilde{u}_t - V(g(\mu_{t-1})) - \delta u \tag{1.5.77}$$

Staying is optimal iff

$$\tilde{u}_t - V(g(\mu_{t-1})) - \delta u > -u \tag{1.5.78}$$

$$\iff \tilde{u}_t > V(g(\mu_{t-1})) - (1 - \delta)u \tag{1.5.79}$$

Which cannot hold since $V(g(\mu_{t-1})) - (1 - \delta)u = \hat{u}_t > \tilde{u}_t$. Therefore, staying in the network is not a profitable deviation for user i .

1.5.10 Proof of Lemma 4

Revenue from decentralizing at size μ is equal to

$$\sum_{t=0}^T \delta^t \frac{1}{8} \frac{\phi(g^{(t)}(\mu_t))^2}{k g^{(t)}(\mu_t)} \quad (1.5.80)$$

and can be approximated by

$$\frac{1}{1-\delta} \left(\frac{1}{8} \frac{\phi(\mu_t)^2}{k \mu_t} + \epsilon(\mu_t) \right) \quad (1.5.81)$$

and given that $g(\mu_t) - \mu_t \rightarrow 0$ as $\mu_t \rightarrow \infty$ it holds that $\epsilon(\mu_t) \rightarrow 0$ as $\mu_t \rightarrow \infty$. Revenue from staying centralized and exploiting (given that μ_t is large enough) is equal to

$$\frac{1}{1-\delta} \psi(\mu_t, -(1-\delta)u) \quad (1.5.82)$$

Thus, it is not sequentially optimal to remain centralized if

$$\frac{1}{1-\delta} \psi(\mu_t, -(1-\delta)u) > \frac{1}{1-\delta} \left(\frac{1}{8} \frac{\phi(\mu_t)^2}{k \mu_t} + \epsilon(\mu_t) \right) \quad (1.5.83)$$

$$\iff \psi(\mu_t, -(1-\delta)u) - \frac{1}{8} \frac{\phi(\mu_t)^2}{k \mu_t} > \epsilon(\mu_t) \quad (1.5.84)$$

Now, I show that the LHS of this inequality is positive and increasing. Once that has been shown, the inequality follows, since the RHS of the inequality goes to 0 for μ_t large enough. To establish that the LHS of the inequality is positive and increasing, consider both possible cases for $\psi(\mu_t, -(1-\delta)u)$. First, consider

$$\mu_t \left(V(\mu_t) + \frac{\phi(\mu_t)^2}{4k\mu_t^2} + (1-\delta)u \right) - \frac{1}{8} \frac{\phi(\mu_t)^2}{k \mu_t} = \mu_t \left(V(\mu_t) + \frac{\phi(\mu_t)^2}{8k\mu_t^2} + (1-\delta)u \right) \quad (1.5.85)$$

Which is both positive and increasing. Next, consider the second case:

$$\sqrt{\frac{V(\mu_t) + (1-\delta)u}{k}} \phi(\mu_t) - \frac{1}{8} \frac{\phi(\mu_t)^2}{k \mu_t} = \phi(\mu_t) \left(\sqrt{\frac{V(\mu_t) + (1-\delta)u}{k}} - \frac{1}{8} \frac{\phi(\mu_t)}{k \mu_t} \right) \quad (1.5.86)$$

Note that to be in the square root part of $\psi(\cdot)$ it holds that $\frac{\phi(\mu_t)}{2k\mu_t} < \sqrt{\frac{V(\mu_t) + (1-\delta)u}{k}}$, which implies that the RHS of the equation above is positive and increasing.

1.5.11 Proof of proposition 4

I start by proving the first statement of the proposition. Suppose it is optimal for the entrepreneur to decentralize in $t = 0$. Now, specifically consider the entrepreneur's revenue

in $t = 1$ when the network has just become operational. By definition of $\psi(\cdot)$, revenue for the entrepreneur in $t = 1$ is maximal for a user utility level of $\hat{u}_t = 0$. Further, note that the entrepreneur's revenues are decreasing in the amount of user utility and that the user utility in a decentralized implementation is strictly positive. Further, the entrepreneur's revenues in $t = 1$ with the decentralized implementation are necessarily smaller for the same user utility level, than if she had retained control of the monetization, by definition of ψ . Thus, period 1 profits are larger for the entrepreneur if she remains centralized. Note however, that the profits of the remaining periods starting in $t = 2$ are necessarily larger for a decentralized implementation, since it was optimal to decentralize in $t = 0$. Thus, it will be sequentially optimal to decentralize the network in $t = 2$, or in a later period, if it is optimal to delay decentralizing again.

Now, I prove the second statement of the proposition. Suppose that it is barely not optimal for the entrepreneur to decentralize her network in $t = 0$, that is, the present value of centralized revenues exceeds that of decentralized revenues by some small amount $\epsilon > 0$. Now, the argument made in the paragraph above shows that delaying decentralization by 1 period increases the entrepreneur's revenue by the amount with which centralized revenues exceed decentralized revenues in $t = 1$. For ϵ small enough, the present value of revenues when decentralizing in $t = 2$ now necessarily exceed those of staying centralized. Thus, the range of locked-in effects for which decentralization is optimal is increased.

Proof of proposition 5

Consider the equilibrium of section 1.3.2 and the incentive of a user i to deviate from punishing the entrepreneur. For that, compare his utility $-u$ from leaving with the utility of staying

$$p \sum_{t=0}^{\infty} \delta^t \hat{u}_t + (1-p)(\hat{u}_t - V(g(\mu_{t-1})) - \delta u) - \epsilon < -u \quad (1.5.87)$$

$$\Rightarrow p \sum_{t=0}^{\infty} \delta^t (V(g^{(t)}(\mu_0) - (1-\delta)u) + (1-p)(-(1-\delta)u - \delta u) - \epsilon < -u \quad (1.5.88)$$

$$\Rightarrow p \sum_{t=0}^{\infty} \delta^t (V(g^{(t)}(\mu_0) - (1-\delta)u) + (1-p)(-(1-\delta)u - \delta u) - \epsilon < -u \quad (1.5.89)$$

$$\Rightarrow p \sum_{t=0}^{\infty} \delta^t (V(g^{(t)}(\mu_0)) < \epsilon \quad (1.5.90)$$

Which is a contradiction for ϵ small enough, i.e. the user is better off when deviating to staying in the network.

Now consider the equilibrium of the main body of the paper. Consider user i 's utility

when staying. If all other users unexpectedly stay in the network, the discounted utility of user i is strictly less than the value of his outside option, since he receives utility $-(1 - \delta)u - \epsilon$ today and discounted future utility equal to $-\delta u$. Thus the utility of staying is equal to $-u - \epsilon < -u$. If, on the other hand, all other users leave the network and user i remains in the network alone, his utility is less than $-(1 - \delta)u - \delta u - \epsilon < -u$ for any $\epsilon > u$. Therefore, user i prefers to stick to the initial equilibrium, regardless of the level of uncertainty p .

1.5.12 Extension: Pre-commitment to revenue sharing path in decentralized governance

Suppose that the entrepreneur can pre-commit to the full path of revenue sharing for all periods $t = 1, 2, \dots$ at the start of the game in $t = 0$. Now, note that for any pre-committed level of α_t , the user's optimal choice of monetization π_t is derived analogously to the optimal monetization π_t^* for a fixed percentage of revenue sharing, and thus equals

$$\frac{1 - \alpha_t}{2k} \frac{\phi(\mu_t)}{\mu_t} \quad (1.5.91)$$

and that the user's utility level for the period thus is

$$V(\mu_t) + \frac{1}{4k} \left((1 - \alpha_t) \frac{\phi(\mu_t)}{\mu_t} \right)^2 \geq 0 \quad (1.5.92)$$

such that the user's choice of monetization implies that it is always optimal for new users to join. Then the entrepreneur's maximization problem in $t = 0$ is equal to

$$\max_{\{\alpha_t\}_{t=1}^{\infty}} \sum_{t=1}^{\infty} \left(\delta^t \frac{\alpha_t (1 - \alpha_t)}{2k} \frac{\phi(\mu_t)}{\mu_t} \right) \quad (1.5.93)$$

Now, straightforward maximization over the α_t implies that in the optimum $\alpha_t = \alpha = 0.5$ for all $t = 1, 2, \dots$

Chapter 2

Mechanism Design for Unequal Societies¹

BASED ON JOINT WORK WITH CARL-CHRISTIAN GROH²

Abstract: We study optimal mechanisms for a utilitarian designer who seeks to assign multiple units of an indivisible good to a group of agents. The agents have heterogeneous marginal utilities of money, which may naturally arise in environments where agents have different wealth levels or financing conditions. The designer faces constraints on ex ante transfers. We show that the ex post efficient allocation rule is not utilitarian optimal in our setting. In certain situations, it is utilitarian optimal to deterministically assign the good to an agent with a lower willingness to pay. This is because a high willingness to pay may stem from a low marginal utility of money. Moreover, the transfer rule does not only facilitate implementation of the desired social choice function in our setting, but also directly affects social welfare. Finally, we highlight how our mechanism can be implemented as an auction with minimum bids and bidding subsidies.

Keywords: optimal mechanism design, redistribution, inequality, auctions

JEL Classification: D44, D47, D61, D63, D82

2.1 Introduction

Consider the following canonical mechanism design problem with a twist: The designer owns a limited number of indivisible goods and a finite number of ex-ante heterogeneous agents with different marginal utilities of money are vying for the allocation of this good.

¹We thank Vitali Gretschko, Hans-Peter Gruner, Volker Nocke, and Thomas Troger for insightful comments. This work was supported by the University of Mannheim's Graduate School of Economic and Social Sciences. Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (Project B03) and the SFB 884 Political Economy of Reforms is gratefully acknowledged. Declarations of interest: none.

²Department of Economics, University of Mannheim

This optimization problem is faced by many entities in the real world. In frequency auctions or airport slot allocation mechanisms, incumbents with large amounts of collateral have significantly easier financing conditions than potential entrants. In government real estate auctions, corporations compete with private citizens. The assignment of kindergarten spots is conducted by small-scale kindergarten providers in many places and the pool of applicants contains families with significantly different wealth levels.³

Designers in these settings will face constraints on transfers. We consider two such constraints motivated by our examples, namely: (1) the constraint that no agent may receive payments from the mechanism in expectation and (2) the requirement that the designer's budget must be balanced ex ante. In any commercial auction conducted by a member state of the European Union (EU), a violation of constraint (1) would imply that the country grants state aid to a participant, creating the possibility for legal penalties.⁴ In our kindergarten example, the local providers we have in mind generally operate under tight budget constraints - thus, kindergarten spot allocations must satisfy constraint (2).

The preferences of the responsible entities may be well reflected by the utilitarian social welfare function. The goals of frequency auctions in the US are, among others, the "efficient use of the spectrum" and "the rapid deployment of new systems".⁵ Similarly, many kindergarten providers such as municipal entities are non-profit organizations. In an attempt to guide the design of mechanisms in these examples, we study the following questions: What allocation rule should a utilitarian designer choose in the small-market allocation problems we have outlined? How can the optimal allocation rule be implemented by slightly altering widespread mechanisms such as auctions?

We show that the utilitarian optimal allocation rule is not ex-post efficient when the marginal utilities of money are heterogeneous across agents. When all agents have the same marginal utility of money, the utilitarian optimal allocation rule is equivalent to the ex post efficient allocation rule, in which the goods are allocated to the agents that state the highest willingnesses to pay. In that setting, the sole purpose of the transfer rule is to

³In Germany, for instance, kindergarten spot allocation decisions are largely made by municipal entities, parent associations, and non-for profit organizations - as detailed by Fritz (2021). Moreover, local authorities condition their decision on several observables that are linked with income - see LH-Mainz (2018).

⁴As defined in article 107 of the TFEU (1957), a government agency grants state aid as defined by European Law if "an intervention gives the recipient an advantage on a selective basis, for example to specific companies or industry sectors" and when "competition has been or may be distorted". An auction setup where the government systematically siphons money to a certain company would thus represent state aid. By contrast, a payment scheme that transfers no money to any company in expectation would arguably not grant any specific company an advantage nor distort competition, thus complying with EU law.

⁵For details, please see Crippen (2000).

implement the desired social choice function because transfers are neutral in terms of social welfare. If the agents have different marginal utilities of money, this is no longer true. This notion breaks the equivalence between the ex-post efficient allocation rule and the utilitarian optimal allocation rule. It is utilitarian optimal to deterministically allocate the good to agents with lower willingnesses to pay in some states of the world. Moreover, the utilitarian optimal allocation rule leaves some units of the good unallocated in certain states of the world when the designer can redistribute. Both of these features imply that the utilitarian optimal allocation rule is not ex-post efficient.

When agents have different marginal utilities of money, the designer will have a redistributive motive. The role of inequality and redistributive preferences of the designer for the allocation of scarce resources has been studied in several papers that precede our work. Weitzman (1977) analyses when a simple rationing scheme in which all consumers get the same amount of a good is preferable to a market price rule. We go a step further and derive a utilitarian optimal mechanism. Condorelli (2013) provides a methodological contribution that enables the derivation of optimal mechanisms for generalized social welfare functions in small markets. We apply this general methodology, but our two particular setups are not discussed by Condorelli (2013). Our preference framework resembles Dworzak et al. (2021) and Akbarpour et al. (2020). By contrast, these authors study settings with a continuum of goods and agents while we study small markets and consider different research questions. Moreover, the designer's willingness and ability to redistribute is central for the results of Dworzak et al. (2021) and Akbarpour et al. (2020), while we derive a bulk of our results for a situation where the designer is unable to redistribute across agents.

We study the following framework: A utilitarian designer initially owns m units of an indivisible good which can be allocated to $N > m$ agents with unit demand for this good. Following Dworzak et al. (2020), an agent's utility consists of two parts. An agent receives utility v_i^K when being allocated the good. Moreover, agents attain utility from the money that they receive from the mechanism. The marginal utility of money for an agent, which we call v_i^M , is constant for each agent but varies across agents. Both v_i^K and v_i^M are private information, but the joint distribution of these variables is common knowledge. In this framework, we characterize the utilitarian optimal mechanism which obeys individual rationality, Bayesian incentive compatibility⁶, and satisfies either (1) the constraint that no agent may receive payments from the mechanism in expectation or (2)

⁶By the results of Gershkov et al. (2013), there also exists an equivalent dominant-strategy implementable mechanism that yields the same utilitarian social welfare as our Bayesian incentive-compatible mechanism.

ex-ante budget balance.⁷

We derive the optimal mechanism based on the following ideas: First, we note that a sufficient statistic for individual behavior is the willingness to pay, namely $r_i = v_i^K / v_i^M$. Dworzak et al. (2021) show that restricting attention to mechanisms that elicit only r_i is without loss of optimality in the settings we study. Secondly, changes in the transfer rules are not neutral in terms of social welfare when agents have heterogeneous marginal utilities of money. The total effect of allocating a good to an agent with a given type r on social welfare is captured by the key statistic of our model: The *inequality adjusted valuation*. In the utilitarian optimal mechanism, the goods will always be allocated to the agents with the highest positive inequality adjusted valuations, which are not necessarily the agents with the highest willingnesses to pay.

Consider the setting where no agent can receive payments from the mechanism in expectation. Because these no-subsidy constraints rule out ex-ante redistribution of money, a utilitarian designer will give any revenue that is raised from an agent back as an ex-ante transfer. In this setting, allocating the good to an agent impacts the expected utility of this agent, and thus social welfare, in three ways. Firstly, it grants the agent consumption utility. When the marginal utilities of money are heterogeneous, consumption utility is not directly inferable from willingness to pay. Secondly, allocating the good to the agent is associated with transfers from the agent to the designer which ensure that incentive compatibility is satisfied. Such changes in the transfer rule impact social welfare. In addition, they affect the expected revenue that the designer raises from the agent, which changes the ex ante transfer that the agent receives - the third effect. When an agent's marginal utility of money is stochastic, acquiring revenue from this agent when her marginal utility of money is below average to fund ex-ante transfers (i.e. redistributing within a given agent) is beneficial. The inequality adjusted valuation condenses these three notions.

To see why allocation by willingness to pay is not optimal in this framework, suppose that different agents have heterogeneous, but deterministic, marginal utilities of money. In that case, the inequality adjusted valuation of any agent equals her consumption utility v_i^K , because the second and third effects mentioned above cancel each other out. Thus, there will be situations where the good is optimally allocated to agents with the lower willingness to pay (but the higher v_i^K). Why would the designer not assign the good to the agent with the higher willingness to pay and compensate the other agent with a payment somewhere in between the willingnesses to pay of the agents? After all, such a deviation would increase the utility of both agents in that particular state of the world. However,

⁷The fact that we only require the budget to be balanced ex ante and not in every possible state of the world is without loss of generality, given the insights of Borgers and Norman (2009).

committing to not executing this trade will be beneficial overall. Recall that the expected payment any agent receives from the designer must be zero. Thus, any money that is moved between the agents would have to be counterbalanced by the designer (either ex ante or in some other states of the world) to ensure that the expected transfers remain zero for each agent. Because the marginal utilities of money are deterministic, the utility any agent obtains from transfers as such must remain unaffected by the aforementioned changes. However, implementing the trade would imply that the agent with the lower v_i^K ultimately receives the good, which is suboptimal.

Now suppose that the designer just faces the ex-ante budget balance requirement discussed earlier. The fact that ex-ante redistribution from one agent to another is now possible endows the designer with a desire to raise revenue from agents with low marginal utilities of money and redistribute this in the form of ex-ante transfers. This additional motive will be reflected in the inequality adjusted valuations of the agents, which still condense all the notions at play under the no-subsidy constraints. Under the ex-ante budget balance condition, there are situations where some units of the good are optimally left unallocated. Mirroring the insights of Myerson (1981), this is beneficial for social welfare because it raises the revenue that is available for redistribution.

To fix ideas, reconsider the two-agent setting introduced previously. Suppose that agent 1 has the higher expected marginal utility of money. This agent's inequality adjusted valuation is the same as in the previous setting with the no-subsidy constraints. However, the designer is now able to redistribute any revenue he raises from agent 2 to agent 1 ex ante. This is reflected by agent 2's inequality adjusted valuation. Allocating the good to agent 2 changes the expected revenue the designer receives from this agent, and thus agent 1's ex-ante transfer, by the virtual valuation of agent 2.⁸ Thus, allocating the good to agent 2 will enable redistribution when agent 2's virtual valuation is positive.

As before, the goods will be allocated to the agents with the highest inequality adjusted valuations, provided they are positive. There will be situations where the good is allocated to the agent with the lower willingnesses to pay, because the agents' inequality adjusted valuations are not equal to their willingnesses to pay. In these situations, changing the allocation such that the agent with the higher willingness to pay receives the good would reduce utilitarian welfare. Committing to refrain from allocating the good to an agent with the higher willingness to pay, but lower inequality adjusted valuation, will enable the designer to generate more revenue from this agent, which is beneficial. For instance, implementation of the ex-post efficient allocation rule through a second-price auction would

⁸As defined in Myerson (1981), the virtual valuation is the maximal amount of revenue that can be raised from an agent in exchange for the allocation of the good at a given willingness to pay.

entail allocation of the goods at prices below those attained in the utilitarian optimal mechanism.

To guide the practical implementation of our ideas, we provide auction rules which implement the respective utilitarian optimal mechanisms in Bayesian equilibrium. Under the no-subsidy constraints, agents with high marginal utilities of money receive bidding subsidies which allow them to compete against agents with easier financing conditions. Under ex-ante budget balance, a similar result holds and the appropriate auction also features bidder-specific minimum bids. We calculate the optimal bidding subsidies and minimum bids for particular examples in section 5.

The rest of our paper proceeds as follows: We offer a detailed literature review in section 2. In section 3, we outline our framework. Section 4 is devoted to the characterization of the optimal mechanisms in the two aforementioned settings. Afterwards, we discuss the implementation of our mechanisms through auctions in section 5, provide some numerical illustrations in section 6, and conclude thereafter.

2.2 Related Literature

In our setting, the role of transfers goes beyond the implementation of the desired allocation rule. Thus, our work relates to three strands of literature. Firstly, our research has strong connections to the contributions that characterize optimal mechanisms in non-quasilinear settings. Secondly, our work relates to the contributions from various fields which investigate the role of heterogeneous, but constant, marginal utilities of money. Thirdly, some of our key ideas complement insights from social choice theory and public finance.

One of the earliest extensions of the standard quasilinear framework was Maskin and Riley (1984), who pin down the optimal auction in a setting with risk-averse buyers. Saitoh and Serizawa (2008), Hashimoto and Saitoh (2010), and Kazumura et al. (2020) characterize, among others, the set of mechanisms that retain certain desiderata in non-quasilinear settings, such as the VCG features. Eisenhuth (2019) studies the revenue-maximizing auction when agents are loss averse and the reference point is endogenous to the choice of the mechanism. Pai and Vohra (2014) and Kotowski (2020) analyse, among others, allocation problems where buyers face heterogeneous budget constraints.

Within this literature, the paper that is most closely related to our own is Huesmann (2017), who examines the problem of assigning a number of indivisible goods to a unit mass of agents with two different wealth levels. All agents have the same underlying

utility function and have concave utility-for-money. Agents only differ in their wealth levels and an agent's wealth is private information. This setup differs from our own in the following ways: Firstly, Huesmann (2017) assumes that an agent's preferences are fully pinned down by the agent's type report (wealth). In our framework, the preferences of an agent are not fully known by the designer, even conditional on the reported willingness to pay. Secondly, she assumes that the utility an agent receives when consuming the good is identical across agents, which we do not. Finally, Huesmann (2017) models a situation with a continuum of agents, whereas we model a finite number of agents to understand the local allocation problems we have in mind. Our key result, namely that it may be utilitarian optimal to deterministically allocate the good to an agent with the lower willingness to pay in certain situations, is unobtainable in the framework of Huesmann (2017). In addition, we show how a designer may account for wealth inequality by implementing an auction with bidding subsidies and minimum bids.

The second related strand of literature consists of papers that study settings where agents differ in their marginal utilities of money, but there are no wealth effects. Esteban and Ray (2006) study a lobbying framework where different lobby groups have different wealth levels and the costs of lobbying fall in wealth. Kang and Zheng (2019) characterize the set of interim-pareto-optimal mechanisms in a setting where one "good" and one "bad" are to be allocated. In their framework, agents have identical values for the "good" and the "bad", but have different (constant) marginal values of money.

Within this literature, the papers that are closest to our own are Dworzak et al. (2021) and Akbarpour et al. (2020). Our modeling technique, in particular the utility function with two dimensional types, is based on the setup in Dworzak et al. (2021). However, both these papers consider a setting with a continuum of goods to be allocated and a continuum of agents to allocate the goods to. By contrast, we model settings with a finite number of agents and goods to be allocated.⁹ Our two main contributions are the closed form expressions governing the utilitarian optimal allocation decision for any given realization of types in our two settings. By construction, these have no counterpart in Dworzak et al. (2021) and Akbarpour et al. (2020), because such characterizations are not required in these papers due to the large markets assumption. Moreover, we derive a bulk of our results under the constraint that no agent can receive positive transfers from the mechanism in expectation, which is not considered by either of these two papers. In addition, our results regarding the implementation of the utilitarian optimal mechanisms via auctions are exclusive to our paper.

⁹In the examples we have mentioned, the number of goods to be allocated and the number of agents vying for the allocation of the goods are small. In such local markets, feasibility constraints have to hold for every possible type realization.

Moreover, both Dworzak et al. (2021) and Akbarpour et al. (2020) consider research questions that differ from the one we study. Dworzak et al. (2021) pin down optimal trading mechanisms in markets with a distinct buyer and seller side where the designer chooses the mechanism. Akbarpour et al. (2020) investigate under what conditions the optimal mechanism in settings where agents have heterogeneous pareto weights is a market-based mechanism such as an auction or employs random allocation.

Condorelli (2013) outlines a method for determining the optimal allocation of goods under generalized objectives of the planner in small-market situations, subject to incentive compatibility and individual rationality. We apply the linear programming approach outlined by Condorelli (2013) to the two particular allocation problems of our paper, both of which are not discussed in Condorelli (2013). In Condorelli (2013), allocation is based on exogenously given priority functions. The inequality adjusted valuations in our model can be understood as endogenous counterparts of these priority functions. These inequality adjusted valuations are determined through the interplay of the incentive compatibility, the individual rationality, and the transfer constraints in our settings.

The idea of assigning different agents heterogeneous welfare weights based on their economic standing was already voiced by Diamond and Mirrlees (1971) and Atkinson and Stiglitz (1976). Our paper is also related to Weitzman (1977), who analyses when a simple rationing scheme in which all consumers get the same amount of a good is preferable to a market price mechanism. Not surprisingly, the advantage of the price based system is increasing in the heterogeneity of taste for the product and falling in the level of inequality. The idea of using the public provision of goods as a redistributive tool is also reflected in the work of Besley and Coate (1991) and Gahvari and Mattos (2007). The authors study a market for an indivisible and rivalrous good such as healthcare. A state with utilitarian objectives will provide an intermediate quality of the good at no costs, which a redistributive act under lump-sum taxation.

2.3 Framework

We consider a finite but arbitrary number of agents $i \in \{1, 2, \dots, N\}$ with unit demand for an indivisible good. Initially $m < N$ units of this good are owned by the mechanism designer and are to be allocated among the agents. Following Dworzak et al. (2021), the agents' behavior is described by the utility function $u_i = v_i^K x_i^K + v_i^M x_i^M$, where v_i^K represents the valuation for the good (which we sometimes refer to as the consumption utility) and x_i^K is a binary variable that describes whether or not the agent has received the good. What sets this specification apart from most of the literature is that the marginal utility

of money may vary across agents. More precisely, the utility derived from money consists of two parts: it equals the marginal utility of money of the agent, namely v_i^M , multiplied by the amount of money received or paid by the agent in the mechanism, namely x_i^M . Both v_i^K and v_i^M are assumed to be private information. However, the joint distribution of these variables, namely F_i , is common knowledge. The marginal densities of v_i^K and v_i^M are denoted by $f_i^K(\cdot)$ and $f_i^M(\cdot)$, respectively. We assume that the mechanism designer is utilitarian and wants to maximize the ex ante welfare given by

$$\sum_{i=1}^N \mathbb{E}[v_i^K x_i^K + v_i^M x_i^M] \quad (2.3.1)$$

subject to incentive compatibility, individual rationality and potential constraints on the transfer rules. Everything else equal, moving money between the agents thus impacts social welfare. We denote the allocation rule by x_i and the transfer rule by t_i . In line with the standard definitions of the literature we say that a mechanism is (Bayesian) incentive compatible if and only if for all agents i and possible types (v_i^K, v_i^M)

$$\begin{aligned} & \mathbb{E}_{-i}[v_i^K x_i(v_i^K, v_i^M, v_{-i}^K, v_{-i}^M) + v_i^M t_i(v_i^K, v_i^M, v_{-i}^K, v_{-i}^M)] \\ & \geq \mathbb{E}_{-i}[v_i^K x_i(\hat{v}_i^K, \hat{v}_i^M, v_{-i}^K, v_{-i}^M) + v_i^M t_i(\hat{v}_i^K, \hat{v}_i^M, v_{-i}^K, v_{-i}^M)] \end{aligned} \quad (2.3.2)$$

holds for all other possible type reports $(\hat{v}_i^K, \hat{v}_i^M)$. We say that participation in a mechanism is individually rational if and only if for all agents and possible types (v_i^K, v_i^M)

$$\mathbb{E}_{-i}[v_i^K x_i(v_i^K, v_i^M, v_{-i}^K, v_{-i}^M) + v_i^M t_i(v_i^K, v_i^M, v_{-i}^K, v_{-i}^M)] \geq \underline{U}_i \quad (2.3.3)$$

holds true, where \underline{U}_i denotes the utility attached to each agent's outside option. Because utility functions are linear in both components, one can normalize the outside option to 0. The transfer t_i represents the money that the agent receives from or pays to the mechanism. Moreover, v_i^K represents the utility gain achieved when receiving the good.

In section 4.1., we require that the expected transfer any agent receives from the mechanism must be weakly negative, i.e. that the following no-subsidy constraints hold for all agents i :

$$\mathbb{E}[t_i(v_i^K, v_i^M, v_{-i}^K, v_{-i}^M)] \leq 0 \quad (2.3.4)$$

In section 4.2., we restrict attention to mechanisms that satisfy ex ante budget balance. We say that a mechanism satisfies ex ante budget balance if and only if

$$\sum_{i=1}^N \mathbb{E}[t_i(v_i^K, v_i^M, v_{-i}^K, v_{-i}^M)] \leq 0 \quad (2.3.5)$$

Restricting attention to a budget balance condition that is expressed in ex ante and not in ex post terms is without loss, given the insights of Borgers and Norman (2009).¹⁰ Note that we do not require transfers to be negative for any possible realization of types. Before deriving the optimal mechanism, we establish some preliminary results following the approach of Dworzak et al. (2021). Because von Neumann-Morgenstern utility functions are only unique up to affine transformations, an agent's rate of substitution $r_i = v_i^K/v_i^M$ will be sufficient to describe her behavior. As derived in Dworzak et al. (2021), an agent's utility function can be rewritten as follows:

$$\mathbb{E}[v_i^K x_i^K + v_i^M x_i^M] = \mathbb{E}_{r_i} [\underbrace{\mathbb{E}[v_i^M | r_i]}_{\lambda_i(r_i)} (r_i x_i^K + x_i^M)] \quad (2.3.6)$$

We assume that for every agent i the rate of substitution r_i is independently and continuously distributed on an interval $[\underline{r}_i, \bar{r}_i]$ with $0 \leq \underline{r}_i \leq \bar{r}_i$. The cdf of r_i will be denoted by $G_i(r_i)$. This distribution is pinned down by the joint distribution of v_i^K and v_i^M . Further note that the factor $\lambda_i(r_i)$ represents a Pareto weight.

Because the statistic r_i fully pins down an agent's behaviour, any attempt at treating two agents with the same r_i (but potentially heterogeneous realizations of v_i^M and v_i^K) differently can not be successful. Thus, restricting attention to mechanisms that elicit only r_i is without loss of optimality. This intuition is formalized in the following proposition due to Dworzak et al. (2021):

Proposition 1 (Dworzak et al. (2021), Theorem 8) *If a mechanism is feasible (respectively, optimal) in the two dimensional model, then there exists a payoff-equivalent mechanism eliciting only r_i that is feasible (respectively, optimal) in the one dimensional model with G_i being the distribution of r_i under the joint distribution F_i for v_i^K and v_i^M , where λ_i is given by:*

$$\lambda_i(r_i) = \mathbb{E}_i[v_i^M | r_i] \quad (2.3.7)$$

In light of this result, we restrict ourselves to mechanisms that elicit only the rate of substitution r_i . Due to the revelation principle, we are also free to restrict our attention to

¹⁰This is because we study independent types and assume that the agents' marginal utilities of money are constant. Proposition 2 in Borgers and Norman (2009) thus establishes the following: For any ex ante budget balanced mechanism in our framework, there exists an ex post budget balanced mechanism with the same allocation rule and interim expected payments (and thus the same utilitarian welfare).

direct mechanisms subject to incentive compatibility constraints. As derived in Dworzak et al. (2021), characterizing incentive compatibility follows the familiar formulation of the literature. Let $X_i(r_i) = \mathbb{E}_{-i}[x_i(r_i, r_{-i})]$ be the expected allocation probability of agent i , given type report r_i , and let $T_i(r_i) = \mathbb{E}_{-i}[t_i(r_i, r_{-i})]$ be the expected transfer of agent i , given type report r_i . Incentive compatibility is characterized by the following lemma:

Lemma 1 (Incentive Compatibility) *A mechanism $\{x_i(r_i, r_{-i}), t_i(r_i, r_{-i})\}_{i=1}^N$ is incentive compatible if and only if*

1. $X_i(r_i)$ is non-decreasing in r_i (Monotonicity)
2. $r_i X_i(r_i) + T_i(r_i) = U_i(\underline{r}_i) + \int_{\underline{r}_i}^{r_i} X_i(s) ds$ (Integrability)

This result follows after rescaling the utility functions and then applying the standard arguments of the literature. By integrability, the expected transfer of an agent is given by

$$\mathbb{E}[T_i(r_i)] = U_i(\underline{r}_i) - \int_{\underline{r}_i}^{\bar{r}_i} X_i(r_i) \underbrace{J_i(r_i)}_{r - \frac{1-G_i(r_i)}{g_i(r_i)}} dG_i(r_i) \quad (2.3.8)$$

where $J_i(r_i)$ denotes the virtual valuation of an agent as defined in Myerson (1981). Note that the integrability condition implies that participation of an agent in the mechanism is individually rational if and only if it is individually rational for the lowest type of the agent. We can use the integrability condition, together with a change of the order of integration, to rewrite the ex ante utilitarian welfare as follows:

$$\sum_i \mathbb{E}[\lambda_i(r_i)(r_i X_i(r_i) + T_i(r_i))] = \sum_i \left(\Lambda_i U_i(\underline{r}_i) + \int_{\underline{r}_i}^{\bar{r}_i} X_i(s) \underbrace{\frac{\int_s^{\bar{r}_i} \lambda_i(r_i) dG_i(r_i)}{g_i(s)}}_{\Pi_i(s)} dG_i(s) \right) \quad (2.3.9)$$

We remark that we have defined $\Lambda_i := \mathbb{E}[\lambda_i(r_i)] = E[v_i^M]$ and:

$$\Pi_i(s) := \frac{\int_s^{\bar{r}_i} \lambda_i(r_i) dG_i(r_i)}{g_i(s)} \quad (2.3.10)$$

Note that for $\lambda_i(r_i) = 1$, we have $\Pi_i(r_i) = \frac{1-G_i(r_i)}{g_i(r_i)}$, the standard inverse hazard rate formulation. Therefore, it seems instructive to think about the function $\Pi_i(r_i)$ as an inequality adjusted inverse hazard rate. Before moving forward, we make the following assumption for the remainder of the paper:

Assumption 1 *For all agents, there exists an $\hat{r}_i \in [\underline{r}_i, \bar{r}_i]$ such that:*

- For all $r_i < \hat{r}_i$, it holds that $J_i(r_i) \leq 0$
- For all $r_i \geq \hat{r}_i$, it holds that $J_i(r_i) \geq 0$

Assumption 1 means that the virtual valuation can cross 0 at most once. Moreover, monotonicity of $J_i(r_i)$ implies assumption 1 but not vice versa. This assumption is useful because, together with monotonicity of $X_i(r_i)$, it is sufficient to ensure that the expected revenue that is raised from an agent in exchange for allocation of the goods will always be weakly positive. A violation of this property in the optimal mechanism would be quite unintuitive, because the agents are generally willing to pay for the consumption of the good. The regularity condition laid out in assumption 1 rules out such outcomes.

2.4 Optimal mechanisms

2.4.1 No-subsidy constraints

For this subsection, we assume that the expected transfer any agent receives from the mechanism must be weakly negative. These constraints make it impossible for the planner to provide ex ante (positive) transfers to any agent. Plugging in the integrability condition, these constraints may be expressed as follows:

$$\mathbb{E}[T_i(r_i)] = U_i(\underline{r}_i) - \int_{\underline{r}_i}^{\bar{r}_i} X_i(r_i) J_i(r_i) dG_i(r_i) \leq 0 \quad \forall i \quad (2.4.1)$$

Noting the way in which we have rewritten the utilitarian social welfare function, the maximization problem can be stated as:

$$\begin{aligned} & \max_{\{x_i(r_i, r_{-i}), U_i(\underline{r}_i)\}_{i=1}^N} \sum_i \left(\Lambda_i U_i(\underline{r}_i) + \int \Pi_i(r_i) x_i(r_i, r_{-i}) dG(r_i, r_{-i}) \right) \\ & \text{s.t. } U_i(\underline{r}_i) - \int_{\underline{r}_i}^{\bar{r}_i} X_i(r_i) J_i(r_i) dG_i(r_i) \leq 0 \quad \forall i \quad (\text{No-subsidy}) \\ & \quad 0 \leq x_i(r_i, r_{-i}) \leq 1 \quad (\text{Prob}) \\ & \quad \sum_i x_i(r_i, r_{-i}) \leq m \quad (\text{Feas}) \\ & \quad X_i(r_i) \text{ non-decreasing} \quad \forall i \quad (\text{Mono}) \\ & \quad U_i(\underline{r}_i) \geq 0 \quad \forall i \quad (\text{IR}) \end{aligned}$$

Then, the optimization problem boils down to choosing the optimal allocation rule and the optimal utility levels for the lowest type of each agent. In the optimal solution to the above problem, all no-subsidy constraints must bind. Suppose, for a contradiction, that the no-subsidy constraint is slack for some agent i in the optimal solution. Then, the designer

could increase $U_i(\underline{r}_i)$ in compliance with this constraint. This change would not violate any other constraint and would raise social welfare, a contradiction. Plugging in these results into our objective function implies that our problem involves the maximization of the following functional, subject to the remaining constraints:

$$\max_{\{x_i(r_i, r_{-i}), U_i(\underline{r}_i)\}_{i=1}^N} \sum_i \left(\int (\Pi_i(r_i) + \Lambda_i J_i(r_i)) x_i(r_i, r_{-i}) dG(r_i, r_{-i}) \right)$$

Key components of this objective functional are the functions $\Pi_i(r_i) + \Lambda_i J_i(r_i)$, which we label now:

Definition 1 (Inequality adjusted valuation - I) *We define the expression $\gamma_i(r_i) := \Pi_i(r_i) + \Lambda_i J_i(r_i)$ to be the inequality adjusted valuation of agent i under the no-subsidy constraints.*

The optimal mechanism, which revolves around these inequality adjusted valuations, is characterized by the following proposition:

Proposition 2 (Optimal Mechanism - I) *Suppose that γ_i is weakly increasing. When ex ante transfers must be weakly negative, the optimal mechanism assigns the goods to the m agents with the highest $\gamma_i(r_i)$. All units of the good are always allocated.*

To understand the result, consider the following relaxed version of our optimization problem, where the monotonicity and the IR constraints are ignored:

$$\begin{aligned} \max_{\{x_i(r_i, r_{-i}), U_i(\underline{r}_i)\}_{i=1}^N} & \sum_i \left(\int (\Pi_i(r_i) + \Lambda_i J_i(r_i)) x_i(r_i, r_{-i}) dG(r_i, r_{-i}) \right) \\ \text{s.t.} & \quad 0 \leq x_i(r_i, r_{-i}) \leq 1 && \text{(Prob)} \\ & \quad \sum_i x_i(r_i, r_{-i}) \leq m && \text{(Feas)} \end{aligned}$$

The structure of the relaxed problem defined above implies that its solution, which can be found using the linear programming approach outlined by Condorelli (2013), must have a bang-bang property. In the solution to this relaxed problem, the agents with the highest positive inequality adjusted valuations $\gamma_i(r_i)$ should always receive the goods. The assumption that $\gamma_i(r_i)$ is weakly increasing implies that these functions will always be positive. Moreover, said assumption also guarantees that the monotonicity constraint will be satisfied in the solution of the relaxed problem. Finally, it remains to show that all IR constraints are also satisfied in this solution. Because the no-subsidy constraints bind, the utility of an agent i with type \underline{r}_i is given by:

$$U_i(\underline{r}_i) = \int_{\underline{r}_i}^{\bar{r}_i} X_i(r_i) J_i(r_i) dG_i(r_i) \quad (2.4.2)$$

Assumption 1 guarantees that the right-hand side of this expression is always strictly positive when the allocation rule $X_i(r_i)$ satisfies monotonicity, which we have shown to hold true. Thus, all IR constraints will be satisfied and the solution of the relaxed problem also constitutes a solution to the original problem. To gain further intuition for the determinants of the optimal mechanism, consider the following decomposition of $\gamma_i(r)$:

$$\gamma_i(r_i) = \Pi_i(r_i) + \Lambda_i J_i(r_i) \quad (2.4.3)$$

$$= \underbrace{\Lambda_i r_i}_{\text{Efficient allocation}} + \underbrace{\frac{\int_{r_i}^{\bar{r}_i} (\lambda_i(s) - \Lambda_i) dG_i(s)}{g_i(r_i)}}_{\text{Ex interim uncertainty}} \quad (2.4.4)$$

The inequality adjusted valuation captures the total effect of allocating a good to an agent i with type r_i on this agent's expected utility, and thus social welfare. They represent endogenous counterparts of the exogenously given priority functions found in Condorelli (2013). In the standard case when $\lambda_i(r_i) = 1$ holds for any agent i and any r_i , the inequality adjusted valuation equals r_i for all agents.

Suppose firstly that the marginal utility of money of any agent i , namely v_i^M , is non-stochastic. Then, the second term of the function $\gamma_i(r_i)$ becomes zero and the inequality adjusted valuation of any agent is exactly equal to her consumption utility v_i^K . This holds due to the following reasoning: Given that v_i^M is non-stochastic, the total expected utility an agent derives from transfers is pinned down by the expected transfer the agent receives. Since this is fixed at 0 by the no-subsidy constraints, allocating the good to agent i only impacts this agent's expected utility via the consumption utility it generates. Any shifts in the transfer schedule induced by the allocation decisions will not affect this agent's expected utility. Thus, the total benefit of allocating the good to agent i with type r_i is given by v_i^K .

Now suppose that an agent's marginal utility of money is stochastic. Then, an allocation decision affects an agent's expected utility even beyond generating consumption utility. This is because an allocation choice affects the transfer schedule of an agent and these shifts are not neutral when the agent's marginal utility of money is stochastic. Intuitively, this stochasticity endows the designer with a desire to transfer money to the agent when this agent's inferred marginal utility of money, namely $\lambda_i(r_i)$, is high and vice versa. This desire is captured by the second component of the inequality adjusted valuation.

The results of this section show that it is not generally utilitarian optimal to allocate goods to the agents with the highest willingnesses to pay. A high willingness to pay does

not necessarily imply a high consumption utility, in particular for agents with a high expected marginal utility of money. Moreover, stochasticity of the marginal utilities of money implies that the impact of allocation choices on transfer rules must be taken into account when maximizing social welfare.

2.4.2 Ex ante budget balance

In this section, we replace the no-subsidy constraints with an ex ante budget balance condition. The implied optimization problem can be stated as:

$$\begin{aligned}
& \max_{\{x_i(r_i, r_{-i}), U_i(\underline{r}_i)\}_{i=1}^N} \sum_i \left(\Lambda_i U_i(\underline{r}_i) + \int \Pi_i(r_i) x_i(r_i, r_{-i}) dG(r_i, r_{-i}) \right) \\
& \text{s.t. } \sum_i \left(U_i(\underline{r}_i) - \int J_i(r_i) x_i(r_i, r_{-i}) dG(r_i, r_{-i}) \right) \leq 0 \quad (\text{Budget}) \\
& \quad 0 \leq x_i(r_i, r_{-i}) \leq 1 \quad (\text{Prob}) \\
& \quad \sum_i x_i(r_i, r_{-i}) \leq m \quad (\text{Feas}) \\
& \quad X_i(r_i) \text{ non-decreasing} \quad (\text{Mono}) \\
& \quad U_i(\underline{r}_i) \geq 0 \quad (\text{IR})
\end{aligned}$$

We define $\Lambda^* = \max\{\Lambda_i\}$ and set $i^* \in \arg \max_i \Lambda_i$. Note firstly that the ex ante budget balance requirement must bind in any optimal mechanism. Otherwise, $U_{i^*}(\underline{r}_i)$ could be increased, leading to a rise in social welfare. Moreover, the IR constraints of all agents $i \notin \arg \max_i \Lambda_i$ must also bind in the optimal mechanism. If any such constraint were slack, the designer could decrease $U_i(\underline{r}_i)$ for some agent $i \notin \arg \max_i \Lambda_i$ to raise $U_{i^*}(\underline{r}_{i^*})$ by the same amount. This change would satisfy all constraints and improve social welfare because $\Lambda^* > \Lambda_i$. Thus, $U_i(\underline{r}_i) = 0$ must hold for all agents $i \notin \arg \max_i \Lambda_i$. Taken together, these two arguments imply that the budget constraint can be rewritten as follows:

$$U_{i^*}(\underline{r}_{i^*}) = \sum_i \int J_i(r_i) x_i(r_i, r_{-i}) dG(r_i, r_{-i})$$

Plugging these results into the objective function implies that our maximization problem, when ignoring the IR and the monotonicity constraints, becomes the following:

$$\begin{aligned}
& \max_{\{x_i(r_i, r_{-i}), U_i(\underline{r}_i)\}_{i=1}^N} \sum_i \left(\int (\Pi_i(r_i) + \Lambda^* J_i(r_i)) x_i(r_i, r_{-i}) dG(r_i, r_{-i}) \right) \\
& \text{s.t. } 0 \leq x_i(r_i, r_{-i}) \leq 1 \quad (\text{Prob}) \\
& \quad \sum_i x_i(r_i, r_{-i}) \leq m \quad (\text{Feas})
\end{aligned}$$

The structure of this relaxed problem is almost the same as in the previous section, with

the only difference being that the arguments in the objective function are now slightly different. To that end, we define the *inequality adjusted valuation* for this particular setting now:

Definition 2 (Inequality adjusted valuation - II) *We define the expression*

$$\varphi_i(r_i) := \Pi_i(r_i) + \Lambda^* J_i(r_i) \quad (2.4.5)$$

to be the inequality adjusted valuation of agent i under the ex ante budget balance requirement.

Note that the function φ_i in this setting differs from the inequality adjusted valuation of the previous setting (γ_i) only in the factor with which an agent's virtual valuation $J_i(r_i)$ is multiplied. Previously, this was Λ_i , and now it is Λ^* . This reflects the following logic: In the previous setting, any money that was raised from an agent i was refunded to this agent ex ante (raising social welfare by Λ_i), while any such money will now be optimally redistributed to the agent with the highest Λ_i , raising social welfare by Λ^* .

In the standard case when $\lambda_i(r) = 1$, the inequality adjusted valuation $\varphi_i(r_i)$ equals r_i as before. Having established this, we characterize the optimal mechanism when φ_i are weakly increasing functions for all agents i .

Proposition 3 (Optimal Mechanism - II) *Suppose that φ_i is weakly increasing for all agents i . Then, the optimal mechanism assigns the good to the m agents with the highest inequality adjusted valuations $\varphi_i(r_i) = \Pi_i(r_i) + \Lambda^* J_i(r_i)$, given that they are positive.*

Consider a realization of types (r_1, \dots, r_N) for which the number of positive $\varphi_i(r_i)$'s is below m . Then, some units of the good will remain unallocated.

Once more, the structure of the relaxed problem implies that its solution has a bang-bang property and assigns the goods to the agents with the highest inequality adjusted valuations $\varphi_i(r_i)$, provided they are positive. The assumption that $\varphi_i(r_i)$ are all weakly increasing guarantees that the solution to the relaxed problem satisfies the monotonicity constraints. Moreover, assumption 1 (together with monotonicity of $X_i(r_i)$ for all agents) implies that the IR constraint of any agent i^* will be satisfied as well, because this agent is guaranteed positive ex-ante transfers. Once more, the solution to the relaxed problem thus constitutes a solution to the general problem.

We now provide sufficient conditions under which assumption 1 holds and the inequality adjusted valuations we have studied (γ_i and φ_i) are weakly increasing:

Remark 1 Consider an agent i and denote the support of v_i^M for this agent by $[v_i^M, \bar{v}_i^M]$. Suppose that the following holds true for an agent i : (i) $\frac{\partial J_i(r_i)}{\partial r_i} \geq 0$, (ii) $\lambda_i(r_i)$ is weakly decreasing in r_i for all $r_i \in [r_i, \bar{r}_i]$, and (iii) $\bar{v}_i^M \leq 2v_i^M$. Then, both γ_i and φ_i are weakly increasing. Moreover, if point (i) holds true for all agents, assumption 1 is satisfied.

Remark 1 states that monotonicity of the inequality adjusted valuations is satisfied in an environment with three characteristics. First, the virtual valuation must be weakly increasing as in Myerson (1981). Secondly, a high willingness to pay is most likely to be supported, ceteris paribus, by a relatively low expected valuation for money. As Dworzak et al. (2021) point out, this assumption is "fairly natural: Generating an increasing $\lambda_i(r_i)$ would require a very strong positive correlation between v_i^K and v_i^M ." Thirdly, the spread of possible marginal utilities of money for any given agent must be sufficiently small. Intuitively, the last requirement requires that the designer knows an agent's marginal utility of money with relatively high precision and that the remaining stochasticity plays no major role. When the inequality adjusted valuations are non-monotonic, the optimal mechanism can be derived using the ironing procedure put forth by Condorelli (2013), which is based on the classical approach of Myerson (1981).

To build further intuition for our results, we present a decomposition of the inequality adjusted valuation φ_i into three components that highlight the trade-offs the designer faces:

$$\varphi_i(r_i) = \Pi_i(r_i) + \Lambda^* J_i(r_i) \quad (2.4.6)$$

$$= \underbrace{\Lambda_i r_i}_{\text{Efficient allocation}} + \underbrace{(\Lambda^* - \Lambda_i) J_i(r_i)}_{\text{Ex ante transfers}} + \underbrace{\frac{\int_{r_i}^{\bar{r}_i} (\lambda_i(s) - \Lambda_i) dG_i(s)}{g_i(r_i)}}_{\text{Ex interim uncertainty}} \quad (2.4.7)$$

The first and third components are present in the inequality adjusted valuation of the previous section as well - only the second term is new. This second component is added because ex-ante redistribution is now possible, which impacts the effect of allocation choices on social welfare. Intuitively, raising revenue as captured by the virtual valuation $J_i(r_i)$ is beneficial for the designer. This revenue will be redistributed ex ante to the poorest agent, increasing welfare by Λ^* at the cost of the ex ante marginal utility of money of the agent from which the revenue was generated, namely Λ_i . The total effect of this ex-ante movement of money on social welfare is captured by the second term.

Naturally, it is of interest to investigate when our allocation rule simplifies to the ex post efficient allocation rule which allocates the good to the m agents with the highest valuations r_i . As pointed out before, our model simplifies to the standard framework when $\lambda_i(r) = 1$ for all agents and thus yields the ex post efficient allocation rule in that

case. More interestingly, our allocation rule also simplifies to the standard allocation in the well studied i.i.d. environment, as is shown in the next corollary:

Corollary 3 *Suppose that either of the following conditions is met:*

1. $\lambda_i(r_i) = 1$ for all i and r_i
2. The pair (v_i^K, v_i^M) is i.i.d. for all agents i and $\varphi_i(r_i)$ is strictly increasing.

Then, the utilitarian optimal allocation rule is equal to the ex-post efficient allocation rule, i.e. the good is assigned to the m agents with the highest valuations r_i .

Suppose that (v_i^K, v_i^M) is i.i.d. among all agents. From an ex ante standpoint, the mapping from willingness to pay into consumption utility is thus the same for all agents. Moreover, we recall from the discussion of proposition 3 that redistribution in the optimal mechanism is implemented through ex ante redistribution to the agent with the highest expected marginal utility of money. However, if every agent is considered equally rich or poor ex ante, then the designer finds himself unwilling and unable to engage in any kind of redistribution. For these reasons, the designer applies the standard allocation rule.

We have noted that some units of the good may be left unallocated in particular situations when the designer faces the ex ante budget balance condition. We say that rationing occurs whenever some units of the good are left unallocated. Rationing is a key source of allocative inefficiency in our model and hence plays an important role for social welfare. Thus, it is instructive to understand how wealth inequality affects the incidence of rationing. Intuitively, rationing is a part of the optimal solution under the ex ante budget balance requirement because, as in Myerson (1981), it positively impacts the amount of money a designer can raise.

In the optimal mechanism, rationing will occur if more than $N - m$ agents have an inequality adjusted valuation that is negative. There is a particularly interesting connection between the ex ante expected valuation of money of each agent, namely Λ_i , and whether or not an agent may be subject to rationing. We say that an agent is subject to rationing when some units of the good are not allocated but this agent still has demand for the good.

Proposition 4 (Rationing) *Let i^* denote the index of the agent with $\Lambda_{i^*} = \Lambda^*$. Then*

1. *Agent i^* is never subject to rationing*
2. *All agents $i \neq i^*$ may be subject to rationing*

Let $\underline{r}_i = 0$ for all i . Then, all agents except agent i^ will be subject to rationing.*

To understand the quantitative magnitude of rationing in a given setting, we consider the probability that rationing occurs, i.e. the fraction of possible type realizations for which rationing would occur. To fix ideas, we assume that the valuations for money are fixed for any agent, but can vary across agents. This allows us to obtain the following results:

Proposition 5 (Inequality and the probability of rationing) *Assume that the marginal utility of money of all agents is non stochastic. Then, it holds that:*

1. $\frac{\partial \Pr(\varphi_i(r_i) < 0)}{\partial \Lambda^*} \geq 0$ holds for all agents $i \neq i^*$. Thus, when Λ^* increases, the probability with which rationing occurs weakly increases.
2. When $\frac{\partial \Pr(\varphi_i(r_i) < 0)}{\partial \Lambda_i} < 0$, a decrease of $\Lambda_i \neq \Lambda^*$ will imply an increase of the probability with which rationing will occur. Note that $\frac{\partial \Pr(\varphi_i(r_i) < 0)}{\partial \Lambda_i} < 0$ holds true if the virtual valuation is weakly increasing.

Note that a higher probability of rationing reflects a greater extent of allocative inefficiency. Modelling the effects of an increase in wealth inequality offers several degrees of freedom. In general, the effect of such a change on the probability with which rationing occurs depends on how an increase in inequality is modelled. Proposition 5 allows us to make the following definitive statements: When the wealth of the poorest members of society decreases, which is reflected by an increase in Λ^* , the probability with which rationing occurs increases, *ceteris paribus*. Result 2 yields insights into the effects of an increase in wealth inequality along the lines of the development of real wages of men in the USA over the years 1990-2010. In these years, real wages of men have stagnated at the 10th percentile and 50th percentile, while they have gone up by around 1% (annualized) at the 90th percentile - see Donovan and Bradley (2019). Thus, over this period, the real wages of the 90th percentile have risen by 22%. Within our model, this can be viewed as a decrease of Λ_i for the wealthier members of the distribution, while all other Λ_i 's are left unchanged. Result 2 shows that the probability with which rationing occurs will increase as a result of these developments.

2.5 Implementation via auctions

2.5.1 No subsidy constraints

In this subsection, we describe how our optimal mechanism in the presence of the no-subsidy constraints can be implemented as an auction. Assume that there is just one good to be allocated and that r_i is continuously distributed on $[\underline{r}_i, \bar{r}_i]$, with $\underline{r}_i = 0$. Assume further that $\gamma_i(r_i)$ is strictly increasing and define the support of possible inequality adjusted valuations of an agent i as $[\underline{\gamma}_i, \bar{\gamma}_i]$.

In the auction we describe, agents only make payments if they win the auction - the winning agent pays their bid, as in a first-price auction. Unlike in a classical first-price auction, our auction employs bidding subsidies. Thus, the bidder who bids the highest amount will not necessarily win the auction. To describe which agent will win the auction, we define the following functions:

$$\Gamma_i(r_i) = Pr\{\max_{j \neq i} \{\gamma_j(r_j)\} \leq \gamma_i(r_i) | r_i\} \quad (2.5.1)$$

$$\tau_i(r_i) := \begin{cases} r_i - (1/\Gamma_i(r_i)) [\int_{r_i}^{r_i} \Gamma_i(s) ds] & r_i > 0 \\ 0 & r_i = 0 \end{cases} \quad (2.5.2)$$

Note that the functions $\Gamma_i(r_i)$ represent the interim allocation probabilities under the utilitarian optimal allocation rule, which the auction needs to implement in Bayesian equilibrium. Order the N agents in a way that $\bar{\gamma}_1 \leq \dots \leq \bar{\gamma}_{N-1} \leq \bar{\gamma}_N$. Let $b = (b_1, \dots, b_N)$ denote the vector of the agents' bids and consider the following functions:

$$\tilde{\tau}_i(b_i) = \begin{cases} \tau_i^{-1}(b_i) & b_i \leq \tau_i(\bar{r}_i) \\ \bar{r}_i & b_i > \tau_i(\bar{r}_i) \end{cases} \quad (2.5.3)$$

Whenever an agent $i \in \{1, \dots, N-1\}$ places a bid above $\tau_i(\bar{r}_i)$, the function $\tilde{\tau}_i(b_i)$ will map this bid into the type \bar{r}_i .

Now consider agent N . Define \tilde{r}_N as the type of this agent that solves $\gamma_N(\tilde{r}_N) = \bar{\gamma}_{N-1}$. In the utilitarian optimal mechanism, this agent will receive the good with probability 1 if her type r_N is above \tilde{r}_N , which means that $\tau_N(r_N)$ will be flat when $r_N \geq \tilde{r}_N$. This distinguishes agent N from all the other agents, whose probabilities of receiving the good are almost surely strictly below 1. Thus, a slight adjustment of $\tilde{\tau}_N(b_N)$ is necessary to accomodate settings where $\bar{\gamma}_{N-1} < \bar{\gamma}_N$ (and hence $\tilde{r}_N < \bar{r}_N$), and $\tilde{\tau}_N(b_N)$ is defined as follows¹¹:

$$\tilde{\tau}_N(b_N) = \begin{cases} \tau_N^{-1}(b_N) & b_N \leq \tau_N(\tilde{r}_N) \\ \tilde{r}_N & b_N > \tau_N(\tilde{r}_N) \end{cases} \quad (2.5.4)$$

Having defined the functions $\tilde{\tau}_i(b_i)$, we close the definition of the auction rules by specifying that an agent i wins the auction if and only if:

$$\gamma_i(\tilde{\tau}_i(b_i)) \geq \max_{j \in \{1, 2, \dots, N\}} \{\gamma_j(\tilde{\tau}_j(b_j))\} \quad (2.5.5)$$

¹¹If $\bar{\gamma}_{N-1} = \bar{\gamma}_N$, then $\tilde{r}_N = \bar{r}_N$.

This auction has a Bayes-Nash equilibrium in which the utilitarian optimal allocation rule for this setting is implemented.

Proposition 6 *Suppose that $\gamma_i(r_i)$ is strictly increasing. In the aforementioned auction, the profile of bidding functions $(b_1(r_1), \dots, b_N(r_N)) = (\tau_1(r_1), \dots, \tau_N(r_N))$ constitutes a Bayes-Nash equilibrium in which the utilitarian optimal allocation rule is implemented.*

By the results of Milgrom and Segal (2002), these bidding functions constitute an equilibrium because the social choice function they induce satisfies the integrability condition and implements the desired allocation rule.

In practical terms, this auction implements our allocation rule via multiplicative bidding subsidies. To see this, consider a setting with two agents $i \in \{1, 2\}$ that have deterministic marginal utilities of money and $v_i^K \sim U[0, 1]$. Define Λ_1 as agent 1's marginal utility of money and Λ_2 as agent 2's marginal utility of money. In the outlined auction, each agent will bid according to the rule $\tau_i(r_i) = 0.5r_i$. For a given vector of bids (b_1, b_2) , agent 1 will receive the good if and only if¹²:

$$\gamma_1(\tau_1^{-1}(b_1)) \geq \gamma_2(\tau_2^{-1}(b_2)) \iff (\Lambda_1/\Lambda_2)b_1 \geq b_2 \quad (2.5.6)$$

Suppose that agent 1 has tougher financing conditions than agent 2, i.e. $\Lambda_1 > \Lambda_2$. Then, agent 1's bids will be scaled up by a factor greater than 1. Thus, agent 1 will receive bidding subsidies in these auctions. The larger the inequality between the agents, i.e. the higher Λ_1/Λ_2 , the greater will be the bidding subsidies received by the agent 1.

2.5.2 Ex ante budget balance

Now, we move on to describe how the utilitarian optimal mechanism in the presence of an ex-ante budget balance constraint can be implemented as an auction. The environment is the same as before: There is one good to be allocated and all agent's types r_i are continuously distributed on $[0, \bar{r}_i]$. The inequality adjusted valuations of all agents are strictly increasing and the supports of the inequality adjusted valuations are given by $[\underline{\varphi}_i, \bar{\varphi}_i]$, where the agents are ordered such that $\bar{\varphi}_1 \leq \dots \leq \bar{\varphi}_{N-1} \leq \bar{\varphi}_N$.

As before, we consider an auction where an agent pays her bid b_i if and only if she wins the auction. All other agents pay nothing. The auction we will outline features bidding subsidies and minimum bids - thus, the highest bidder does not necessarily win the auction. To specify the assignment rule, we have to specify some auxiliary functions.

¹²The corresponding calculations may be found in appendix 2.8.11.

The utilitarian optimal allocation rule implements the following interim allocation probabilities:

$$\Phi_i(r_i) := Pr\{\max_{j \neq i} \{\varphi_j(r_j)\} \leq \varphi_i(r_i) | r_i\} \quad (2.5.7)$$

Note that the inequality adjusted valuations can be negative in this setting. To that end, define r_i^{min} as the agent's type that satisfies $\varphi_i(r_i^{min}) = 0$. Analogously as before, we define the following functions:

$$\beta_i(r_i) := \begin{cases} r_i - (1/\Phi_i(r_i)) [\int_{r_i}^{r_i} \Phi_i(s) ds] & r_i > r_i^{min} \\ r_i^{min} & r_i \leq r_i^{min} \end{cases} \quad (2.5.8)$$

For all agents $i \leq N - 1$, we define the following functions:

$$\tilde{\beta}_i(b_i) = \begin{cases} r_i^{min} & b_i \leq \beta_i(r_i^{min}) \\ \beta_i^{-1}(b_i) & b_i \in (\beta_i(r_i^{min}), \beta_i(\tilde{r}_i)) \\ \tilde{r}_i & b_i \geq \beta_i(\tilde{r}_i) \end{cases} \quad (2.5.9)$$

For agent N , this function has to be slightly altered to adjust for the fact that this agent will certainly receive the good when $r_N > \tilde{r}_N$, where $\varphi_N(\tilde{r}_N) = \bar{\varphi}_{N-1}$. This function is:

$$\tilde{\beta}_N(b_N) = \begin{cases} r_N^{min} & b_N \leq \beta_N(r_N^{min}) \\ \beta_N^{-1}(b_N) & b_N \in (\beta_N(r_N^{min}), \beta_N(\tilde{r}_N)) \\ \tilde{r}_N & b_N \geq \beta_N(\tilde{r}_N) \end{cases} \quad (2.5.10)$$

Finally, agent i wins the auction if and only if:

$$\varphi_i(\tilde{\beta}_i(b_i)) > 0 \wedge \varphi_i(\tilde{\beta}_i(b_i)) \geq \max_{j \in \{1, 2, \dots, N\}} \{\varphi_j(\tilde{\beta}_j(b_j))\} \quad (2.5.11)$$

The following proposition formalizes that this auction has a Bayes-Nash equilibrium in which our allocation rule is implemented.

Proposition 7 *In the auction described above, the profile of bidding functions $(b_1(r_1), \dots, b_N(r_N)) = (\beta_1(r_1), \dots, \beta_N(r_N))$ constitutes a Bayes-Nash equilibrium, in which the utilitarian optimal allocation rule is implemented.*

In the auction outline above, there are bidder-specific minimum bids. For every agent i , her minimum bid is given by r_i^{min} , since these solve $\gamma_i(r_i^{min}) = 0$. A bidder i has a positive chance of receiving the good if and only if her bid is above the bidder-specific minimum bid r_i^{min} . Note that all agents $i \notin \arg \max_j \Lambda_j$ will have strictly positive minimum bids,

while all agents $i^* \in \arg \max_j \Lambda_j$ will have a minimum bid equal to 0.¹³

To illustrate the above insights, we calculate the auction rules and equilibrium for the following example: Suppose agent 1 has a valuation of the good $v_1^K \sim U[0, 1]$ and a deterministic utility of money $\Lambda_1 = 1$, such that $r_1 \sim U[0, 1]$. Agent 2 has a valuation of the good $v_2^K \sim U[0, 2]$ and a deterministic utility of money $\Lambda_2 = 2$ such that $r_2 \sim U[0, 1]$.

Using equations (2.5.7) and (2.5.8), the equilibrium bidding function of agent 1 equals $b_1(r_1) = 0.5r_1 + \frac{1}{6}$ while agent 2 bids according to $b_2(r_2) = \frac{r_2^2}{2r_2+1}$. Agents 1 and 2 have minimum bids equal to $r_1^{\min} = \frac{1}{3}$ and $r_2^{\min} = 0$, respectively. When agent 1 submits a bid below her minimum bid, she will never win the auction. If agent 1 bids above this minimum bid, she will win the auction if and only if¹⁴:

$$b_1 \geq b_2 + \underbrace{\frac{1}{3} \left(\sqrt{b_2^2 + b_2} - 2b_2 + 1 \right)}_{\text{bidding subsidy}} \quad (2.5.12)$$

This bidding subsidy is non-constant and exactly equal to $\frac{1}{3}$ at bids $b_2 = 0$ and $b_2 = \frac{1}{3}$, and is slightly larger at bids b_2 in between these values.

2.6 Numerical illustrations

To further emphasize the key points of our paper, we provide some numerical illustrations. Assume that there are two agents $i = 1, 2$ with $v_i^K \sim U[0, 1]$ and $v_1^M \sim \text{Pareto}(k = 3, x_{\min} = 1.5)$ while $v_2^M \sim \text{Pareto}(k = 3, x_{\min} = 1)$. Note that the support of the Pareto distribution is $[x_{\min}, \infty)$. Therefore, agent 2 has support on lower values of v^M than agent 1. This can naturally arise in a setting where agent 2 is ex ante more wealthy than agent 1 or has easier financing conditions, but preferences are not fully known ex ante. In this example, both inequality adjusted valuations will be non-decreasing.

In the following figure, we demonstrate how the utilitarian optimal allocation rule is modified in the presence of inequality.¹⁵ The yellow line illustrates the ex post efficient allocation rule. The blue line represents the utilitarian optimal allocation rule under ex ante budget balance. The red line represents the utilitarian optimal allocation rule under the no-subsidy constraints. All three allocation rules can be understood as follows: For

¹³Formally, this holds because the inequality adjusted valuation of any agent i^* is zero when this agent's type is $r_{i^*} = 0$. For any other agent j , the inequality adjusted valuation at $r_j = 0$ is strictly negative.

¹⁴The corresponding calculations may be found in appendix 2.8.11.

¹⁵The algebra involved in calculating the inequality adjusted valuations may be found in appendix A.11.

a given r_1 (on the x-axis), agent 2 is allocated the good under a given allocation rule if and only if her willingness to pay r_2 is such that the point (r_1, r_2) lies above the line corresponding to the allocation rule.




Figure 2.1: Utilitarian optimal allocation rule vs. ex post efficient allocation rule

In either utilitarian optimal mechanism, agent 1 (who is perceived to have more difficult financing conditions) receives the good more often than under the ex post efficient allocation rule. This result is driven by two forces. Firstly, when agent 1 reports low values of r , this is often driven by a high marginal utility of money, not by a low consumption utility. Achieving allocative efficiency necessitates controlling for this. Secondly, the designer can realize his preference to redistribute money from agent 2 to agent 1 under the ex ante budget balance requirement. When agent 2 has a negative virtual valuation, allocating the good to this agent reduces the revenue that is raised from her, which is undesirable to the designer. This second effect leads to greater differences between the utilitarian optimal and the ex post efficient allocation rules when r_2 is low because the virtual valuation is monotonic in this example. Even when $r_1 = r_2 = 0$, agent 1 would surely receive the good under the ex ante budget balance requirement, because allocating the good to agent 2 would negatively impact the revenue that can be raised from the latter.

Finally, it remains to discuss the discrepancies between the utilitarian optimal allocation rules under the no-subsidy constraints and under the ex ante budget balance constraint. Any differences inbetween these rules are exclusively driven by the fact that the designer can redistribute towards agent 2 under ex ante budget balance, but not under the no-subsidy constraints. At low values of r_2 , agent 2's virtual valuation is negative, thus favoring allocation of the good towards agent 1 when the budget has to be balanced (blue line). At high values of r_2 , the opposite holds, thus motivating the assignment of

the good to agent 2. These notions do not affect the designer's choices when he only has to obey the no-subsidy constraints (red line). Taken together, these arguments imply the following: When r_2 is low, there is an interval of r_1 where agent 1 would receive the good when the budget has to be balanced ex ante but not when the designer only faces the no-subsidy constraints. The opposite holds true when r_2 is large. Graphically, this is reflected by the fact that the blue line lies above the red line when r_2 is small and vice versa.

2.7 Conclusion

We have derived the utilitarian optimal mechanism for an assignment problem in which the designer initially owns m units of an indivisible good which are to be distributed among a finite number N of agents, where $N > m$. In contrast to the usual assumption made in the literature, we work with heterogeneous marginal utilities of money. This implies that utility is not perfectly transferable between agents. We have formalized this feature by adapting the model of Dworzak et al. (2021) to our framework. In addition to the standard incentive compatibility and individual rationality constraints, the designer also faces additional constraints on transfers, namely (i) a constraint stating that no agent can receive positive transfers from the mechanism in expectation or (ii) a requirement that the designer's budget must be balanced ex ante. In these settings, it is generally not utilitarian optimal to allocate the goods to the agents with the highest willingnesses to pay.

We derive the utilitarian optimal mechanism using methodologies developed in Condorelli (2013) and Dworzak et al. (2021). It revolves around a key statistic which we call the *inequality adjusted valuation*. The inequality adjusted valuation, whose exact form depends on the constraints on transfers that are imposed, condenses three critical considerations: First, the designer has a desire to allocate the goods to the agents with the highest consumption utilities, ceteris paribus. Second, when there is ex-interim uncertainty about the agent's marginal utilities of money, the designer wants to pay transfers to the agent when her marginal utility of money is above its average and vice versa. The allocation rule will reflect this, because assignment of the good will always be associated with payments. Thirdly, when redistribution is possible under the ex ante budget balance requirement, the allocation rule will affect the revenue that the designer raises for redistribution.

We have shown that the utilitarian optimal mechanism allocates the good to an agent if and only if (i) her inequality adjusted valuation is among the m highest inequality adjusted valuations and (ii) her inequality adjusted valuation is positive. The agents with the

highest inequality adjusted valuations do not necessarily have the highest willingnesses to pay. Thus, heterogeneity in the marginal utilities of money creates a tension between ex-post efficiency and ex-ante optimality. Under the ex ante budget balance condition, there are states of the world in which some units of the good are left unallocated. Such outcomes, which are a byproduct of the designer's revenue motive, may be exacerbated by high levels of inequality.

In addition, we have discussed how the optimal mechanism may be implemented as an auction in which agents with high marginal utilities of money receive bidding subsidies. This mechanism provides a simple way of achieving optimal allocation in the presence of heterogeneity in the marginal utilities of money. Finally, we have illustrated our results with numerical examples and have shown that agents who are perceived to have high marginal utilities of money have a higher chance of receiving the good in either utilitarian optimal mechanism than in the ex post efficient allocation rule.

We have provided further evidence that it is not without loss of generality to normalize the marginal utility of money of all agents to 1 when studying utilitarian social welfare maximization. In the real world applications we have discussed, incorporating our ideas into the implemented mechanisms may be beneficial even beyond raising instantaneous social welfare. In the kindergarden example, accounting for wealth inequality may foster equality of opportunity by promoting equal access to education. In the auction examples we discussed, applying our insights may be quite pro-competitive. This is because our mechanism reduces the advantage that incumbents with easy financing conditions have in traditional auction mechanisms.

2.8 Appendix for Mechanism Design for Unequal Societies

2.8.1 Proof of Proposition 1

See Dworczak et al. (2021).

2.8.2 Proof of Lemma 1

The result follows directly from rescaling the agents' utility functions and applying the standard results from the literature.

2.8.3 Proof of Proposition 2

Part 1: Deriving the relaxed problem

The full optimization problem is given by the following:

$$\begin{aligned}
& \max_{\{x_i(r_i, r_{-i}), U_i(\underline{r}_i)\}_{i=1}^N} \sum_i \left(\Lambda_i U_i(\underline{r}_i) + \int \Pi_i(r_i) x_i(r_i, r_{-i}) dG(r_i, r_{-i}) \right) \\
& \text{s.t. } U_i(\underline{r}_i) - \int J_i(r_i) x_i(r_i, r_{-i}) dG_i(r_i, r_{-i}) \leq 0 \quad \forall i \quad (\text{Transfers}) \\
& \quad 0 \leq x_i(r_i, r_{-i}) \leq 1 \quad (\text{Prob}) \\
& \quad \sum_i x_i(r_i, r_{-i}) \leq m \quad (\text{Feas}) \\
& \quad X_i(r_i) \text{ non-decreasing} \quad (\text{Mono}) \\
& \quad U_i(\underline{r}_i) \geq 0 \quad (\text{IR})
\end{aligned}$$

The transfer constraints must bind for all agents. Otherwise, $U_i(\underline{r}_i)$ could be increased for a given agent i , which is in line with all constraints and would raise welfare.

Plugging in this result into our objective function yields the following:

$$\begin{aligned}
& \sum_i \left(\Lambda_i U_i(\underline{r}_i) + \int \Pi_i(r_i) x_i(r_i, r_{-i}) dG(r_i, r_{-i}) \right) = \\
& \sum_i \left(\Lambda_i \int J_i(r_i) x_i(r_i, r_{-i}) dG_i(r_i, r_{-i}) + \int \Pi_i(r_i) x_i(r_i, r_{-i}) dG(r_i, r_{-i}) \right) = \\
& \sum_i \left(\int (\Pi_i(r_i) + \Lambda_i J_i(r_i)) x_i(r_i, r_{-i}) dG(r_i, r_{-i}) \right) \quad (2.8.1)
\end{aligned}$$

Ignoring the monotonicity and the IR constraints, our optimization problem thus boils down to maximizing the aforementioned function, subject to the feasibility and the probability constraints.

This is a standard linear programming problem. Mirroring the insights of Condorelli (2013), the goods will be allocated to the m agents with the highest $\gamma_i(r_i)$ for any given realization of types in the solution of this relaxed problem.

Part 2: Showing that the solution to the relaxed problem solves the full optimization problem.

It remains to show that both the monotonicity and the IR constraints will be satisfied in the solution to this relaxed problem. Monotonicity will be satisfied because γ_i is increasing in r_i and the allocation probability $X_i(r_i)$ is increasing in γ_i .

Now consider the IR constraint of agent i^* . We show the following: Under assumption 1 and for a monotone $X_i(r_i)$, the integral $\int_{\underline{r}_i}^{\bar{r}_i} J_i(r_i)g_i(r_i)dr_i$ will always be strictly positive.

Firstly, note the following:

$$\begin{aligned} \int_{\underline{r}_i}^{\bar{r}_i} J_i(r_i)g_i(r_i)dr_i &= \left[\int_{\underline{r}_i}^{\bar{r}_i} \left(r_i - \frac{1 - G_i(r_i)}{g_i(r_i)} \right) g_i(r_i) dr_i \right] = \\ &= \left[\int_{\underline{r}_i}^{\bar{r}_i} \left(g_i(r_i)r_i - (1 - G_i(r_i)) \right) dr_i \right] = \left[\int_{\underline{r}_i}^{\bar{r}_i} (g_i(r_i)r_i + G_i(r_i)) dr_i - \int_{\underline{r}_i}^{\bar{r}_i} (1) dr_i \right] \end{aligned} \quad (2.8.2)$$

Note that:

$$\frac{\partial G_i(r_i)r_i}{\partial r_i} = g_i(r_i)r_i + G_i(r_i) \quad (2.8.3)$$

Plugging this in yields that:

$$\int_{\underline{r}_i}^{\bar{r}_i} J_i(r_i)g_i(r_i)dr_i = \left[[G_i(r_i)r_i]_{\underline{r}_i}^{\bar{r}_i} - \int_{\underline{r}_i}^{\bar{r}_i} (1) dr_i \right] = \left[\bar{r}_i - \int_{\underline{r}_i}^{\bar{r}_i} (1) dr_i \right] = \underline{r}_i \quad (2.8.4)$$

Thus, this term is always weakly positive as long as $\underline{r}_i \geq 0$.

By assumption 1, there exists an \hat{r}_i such that:

$$\forall r_i \geq \hat{r}_i : J_i(r_i) \geq 0 \quad (2.8.5)$$

$$\forall r_i < \hat{r}_i : J_i(r_i) \leq 0 \quad (2.8.6)$$

By this assumption and because $g_i(r_i)$ is always positive, it holds that:

$$\forall r_i \geq \hat{r}_i : J_i(r_i)g_i(r_i) \geq 0 \quad (2.8.7)$$

$$\forall r_i < \hat{r}_i : J_i(r_i)g_i(r_i) \leq 0 \quad (2.8.8)$$

Consider a monotone $X_i(r_i)$. By monotonicity of X_i and the above arguments, we have that:

$$\forall r_i \geq \hat{r}_i : X_i(r_i) \geq X_i(\hat{r}_i) \implies X_i(r_i)J_i(r_i)g_i(r_i) \geq X_i(\hat{r}_i)J_i(r_i)g_i(r_i) \quad (2.8.9)$$

$$\forall r_i < \hat{r}_i : X_i(r_i) \leq X_i(\hat{r}_i) \implies X_i(r_i)J_i(r_i)g_i(r_i) \geq X_i(\hat{r}_i)J_i(r_i)g_i(r_i) \quad (2.8.10)$$

Thus, we have that:

$$\begin{aligned} \int_{r_i}^{\bar{r}_i} X_i(r_i)J_i(r_i)g_i(r_i)dr_i &= \int_{r_i}^{\hat{r}_i} X_i(r_i)J_i(r_i)g_i(r_i)dr_i + \int_{\hat{r}_i}^{\bar{r}_i} X_i(r_i)J_i(r_i)g_i(r_i)dr_i \geq \\ \int_{r_i}^{\hat{r}_i} X_i(\hat{r}_i)J_i(r_i)g_i(r_i)dr_i &+ \int_{\hat{r}_i}^{\bar{r}_i} X_i(\hat{r}_i)J_i(r_i)g_i(r_i)dr_i = X_i(\hat{r}_i) \int_{r_i}^{\bar{r}_i} J_i(r_i)g_i(r_i)dr_i = X_i(\hat{r}_i)\underline{r}_i \geq 0 \end{aligned} \quad (2.8.11)$$

The last inequality holds since $X_i(\hat{r}_i)$ is a weakly positive constant.

This result implies that the IR constraints will also be satisfied in the solution of the relaxed problem. Thus, we are done.

Part 3: No rationing

Finally, it remains to show that all units of the good will always be allocated. Sufficient for this is to show that γ_i is always strictly positive.

Thus, we need to show that $\gamma_i(r_i) \geq 0$ for all i . To show this, it suffices to show that $\gamma_i(\underline{r}_i) \geq 0$, together with our assumption that γ_i is increasing in r_i . Thus, note that:

$$\gamma_i(r_i) = \Lambda_i r_i + \frac{\int_{r_i}^{\bar{r}_i} (\lambda_i(s) - \Lambda_i) dG_i(s)}{g_i(r_i)} \quad (2.8.12)$$

$$\implies \gamma_i(\underline{r}_i) = \Lambda_i \underline{r}_i + \frac{\int_{\underline{r}_i}^{\bar{r}_i} (\lambda_i(s) - \Lambda_i) dG_i(s)}{g_i(\underline{r}_i)} = \Lambda_i \underline{r}_i \geq 0 \quad (2.8.13)$$

The last inequality holds because $\Lambda_i = \mathbb{E}[\lambda_i(s)]$.

2.8.4 Proof of Proposition 3

Part 1: Deriving and solving a relaxed problem

The full problem is the following: The optimal mechanism needs to solve:

$$\begin{aligned}
 & \max_{\{x_i(r_i, r_{-i}), U_i(\underline{r}_i)\}_{i=1}^N} \sum_i \left(\Lambda_i U_i(\underline{r}_i) + \int \Pi_i(r_i) x_i(r_i, r_{-i}) dG(r_i, r_{-i}) \right) \\
 & \text{s.t. } \sum_i \left(U_i(\underline{r}_i) - \int J_i(r_i) x_i(r_i, r_{-i}) dG(r_i, r_{-i}) \right) \leq 0 \quad (\text{Budget}) \\
 & \quad 0 \leq x_i(r_i, r_{-i}) \leq 1 \quad (\text{Prob}) \\
 & \quad \sum_i x_i(r_i, r_{-i}) \leq m \quad (\text{Feas}) \\
 & \quad X_i(r_i) \text{ non-decreasing} \quad (\text{Mono}) \\
 & \quad U_i(\underline{r}_i) \geq \underline{U}_i = 0 \quad (\text{IR})
 \end{aligned}$$

Note that the budget constraint must bind in the solution to this problem. Suppose, for a contradiction, that it is slack. Then, $U_i(\underline{r}_i)$ could be increased for some agent i . This would be in line with all constraints and would raise social welfare, implying that the starting mechanism could not have been optimal.

Based on this, note that the IR constraints for all agents $j \notin \arg \max_i \Lambda_i$ must also bind. Suppose, for a contradiction, that there is one such constraint that does not bind. Then, $U_{i^*}(\underline{r}_{i^*})$ could be increased at the cost of a one-for-one decrease in $U_j(\underline{r}_j)$. Because $\Lambda^* > \Lambda_j$ by definition, this would raise welfare without violating any other constraints.

Taking these two results together implies that the budget constraint can be rewritten as follows:

$$\begin{aligned}
 & \sum_i \left(U_i(\underline{r}_i) - \int J_i(r_i) x_i(r_i, r_{-i}) dG(r_i, r_{-i}) \right) = 0 \iff \\
 & \sum_i U_i(\underline{r}_i) = U_{i^*}(\underline{r}_{i^*}) = \sum_i \left(\int J_i(r_i) x_i(r_i, r_{-i}) dG(r_i, r_{-i}) \right) = 0 \quad (2.8.14)
 \end{aligned}$$

Plugging this into the objective function and ignoring the remaining IR constraints and the monotonicity constraints implies that our optimization problem boils down to the

following:

$$\begin{aligned} \max_{\{x_i(r_i, r_{-i}), U_i(r_i)\}_{i=1}^N} & \sum_i \left(\int (\Pi_i(r_i) + \Lambda^* J_i(r_i)) x_i(r_i, r_{-i}) dG(r_i, r_{-i}) \right) \\ \text{s.t. } & 0 \leq x_i(r_i, r_{-i}) \leq 1 && \text{(Prob)} \\ & \sum_i x_i(r_i, r_{-i}) \leq m && \text{(Feas)} \end{aligned}$$

Following the insights of Condorelli (2013), the solution to this relaxed problem assigns the goods to the m agents with the highest φ_i , provided they are positive.

Part 2: Showing that the solution to the relaxed problem also solves the general optimization problem

Finally, it remains to show that the monotonicity constraints and IR constraint of i^* will hold in the solution of this relaxed problem.

Monotonicity requires that $X_i(r_i)$ is non-decreasing. Note our key assumption that $\varphi_i(r_i)$ is increasing in r_i . For values of r_i where the inequality adjusted valuation is negative, monotonicity holds. Now consider values of r_i where $\varphi_i(r_i) > 0$. Agent i will receive the good if and only the agent's inequality adjusted valuation is among the m highest. Since $\varphi_i(r_i)$ is increasing in r_i , this implies that the probability of allocation cannot be falling in r_i .

In the proof of the previous proposition, we have shown that the integrals $\int J_i(r_i) x_i(r_i, r_{-i}) dG(r_i, r_{-i})$ must all be positive, given that $X_i(r_i)$ is monotone. This implies that the IR constraints of all agents i^* must also be satisfied in the solution to the relaxed problem outlined above.

2.8.5 Proof of Remark 1

Part 1: Showing that $\varphi_i(r_i)$ is monotone under the stated assumptions.

Note first that the inequality adjusted valuation φ_i can be written as follows:

$$\varphi_i(r_i) = \Lambda^* J_i(r_i) + \frac{1 - G_i(r_i)}{g_i(r_i)} \mathbb{E}[\lambda_i(s) | s \geq r_i] \quad (2.8.15)$$

Taking the derivative of φ_i w.r.t r_i yields:

$$\frac{\partial \varphi_i}{\partial r_i} = \Lambda^* \frac{\partial J_i(r_i)}{\partial r_i} + \left[\frac{\partial}{\partial r_i} \left(\frac{1 - G_i(r_i)}{g_i(r_i)} \right) \right] \mathbb{E}[\lambda_i(s) | s \geq r_i] +$$

$$\left[\frac{1 - G_i(r_i)}{g_i(r_i)} \right] \frac{\partial \mathbb{E}[\lambda_i(s) | s \geq r_i]}{\partial r_i} \quad (2.8.16)$$

The critical term is the third term. We evaluate this term in more detail. Note firstly that:

$$\mathbb{E}[\lambda_i(s) | s \geq r_i] = \int_{r_i}^{\bar{r}_i} \lambda_i(s) g_i(s) [1 - G_i(r_i)]^{-1} d(s) \quad (2.8.17)$$

The derivative of this w.r.t r_i is:

$$\frac{\partial \mathbb{E}[\lambda_i(s) | s \geq r_i]}{\partial r_i} =$$

$$-\lambda_i(r_i) g_i(r_i) [1 - G_i(r_i)]^{-1} + \int_{r_i}^{\bar{r}_i} \lambda_i(s) g_i(s) [1 - G_i(r_i)]^{-2} (-1) (-g_i(r_i)) d(s) \quad (2.8.18)$$

$$= -\lambda_i(r_i) \left[\frac{1 - G_i(r_i)}{g_i(r_i)} \right]^{-1} + \left[\frac{1 - G_i(r_i)}{g_i(r_i)} \right]^{-1} \int_{r_i}^{\bar{r}_i} \lambda_i(s) g_i(s) [1 - G_i(r_i)]^{-1} d(s) \quad (2.8.19)$$

$$= \left[\frac{1 - G_i(r_i)}{g_i(r_i)} \right]^{-1} \left[\mathbb{E}[\lambda_i(s) | s \geq r_i] - \lambda_i(r_i) \right] \quad (2.8.20)$$

Plugging in then yields that:

$$\frac{\partial \varphi_i}{\partial r_i} = \Lambda^* \frac{\partial J_i(r_i)}{\partial r_i} + \left[\frac{\partial}{\partial r_i} \left(\frac{1 - G_i(r_i)}{g_i(r_i)} \right) \right] \mathbb{E}[\lambda_i(s) | s \geq r_i] + \left[\mathbb{E}[\lambda_i(s) | s \geq r_i] - \lambda_i(r_i) \right] \quad (2.8.21)$$

$$= \Lambda^* \frac{\partial J_i(r_i)}{\partial r_i} + \left[\frac{\partial}{\partial r_i} \left(\frac{1 - G_i(r_i)}{g_i(r_i)} \right) - 1 \right] \mathbb{E}[\lambda_i(s) | s \geq r_i] + \left[2\mathbb{E}[\lambda_i(s) | s \geq r_i] - \lambda_i(r_i) \right] \quad (2.8.22)$$

$$= \frac{\partial J_i(r_i)}{\partial r_i} \left[\Lambda^* - \mathbb{E}[\lambda_i(s) | s \geq r_i] \right] + \left[2\mathbb{E}[\lambda_i(s) | s \geq r_i] - \lambda_i(r_i) \right] \quad (2.8.23)$$

The last equality holds because $\frac{\partial J_i(r_i)}{\partial r_i} = 1 - \frac{\partial}{\partial r_i} \left(\frac{1 - G_i(r_i)}{g_i(r_i)} \right)$. The derivative $\frac{\partial \varphi_i}{\partial r_i}$ is hence weakly positive if and only if:

$$\frac{\partial J_i(r_i)}{\partial r_i} \left[\Lambda^* - \mathbb{E}[\lambda_i(s) | s \geq r_i] \right] \geq \lambda_i(r_i) - 2\mathbb{E}[\lambda_i(s) | s \geq r_i] \quad (2.8.24)$$

Now we show that this inequality holds under the sufficient conditions outlined in the remark.

The assumption that $\lambda_i(s)$ is weakly decreasing in s implies that $\mathbb{E}[\lambda_i(s)|s \geq r_i] \leq \lambda_i(r_i)$ and hence:

$$\frac{\partial \mathbb{E}[\lambda_i(s)|s \geq r_i]}{\partial r_i} = \left[\frac{1 - G_i(r_i)}{g_i(r_i)} \right]^{-1} \left[\mathbb{E}[\lambda_i(s)|s \geq r_i] - \lambda_i(r_i) \right] \leq 0 \quad (2.8.25)$$

In turn, this implies the following for any agent:

$$\mathbb{E}[\lambda_i(s)|s \geq r_i] \leq \mathbb{E}[\lambda_i(s)|s \geq \underline{r}_i] = \Lambda_i \leq \Lambda^* \quad (2.8.26)$$

Thus, the assumption that $\lambda_i(s)$ is weakly decreasing in s , together with the assumption that $J'_i(r_i) \geq 0$, implies that:

$$\frac{\partial J_i(r_i)}{\partial r_i} \left[\Lambda^* - \mathbb{E}[\lambda_i(s)|s \geq r_i] \right] \geq 0 \quad (2.8.27)$$

Finally, note that the support of v_i^M is $[\underline{v}_i^M, \bar{v}_i^M]$. Then, it holds that $\lambda_i(s) \in [\underline{v}_i^M, \bar{v}_i^M]$ and thus $\mathbb{E}[\lambda_i(s)|s \geq r_i] \in [\underline{v}_i^M, \bar{v}_i^M]$. The assumption that $\bar{v}_i^M \leq 2\underline{v}_i^M$ thus implies that:

$$\lambda_i(r_i) - 2\mathbb{E}[\lambda_i(s)|s \geq r_i] \leq \bar{v}_i^M - 2\underline{v}_i^M \leq 0 \quad (2.8.28)$$

Summing up: The three stated conditions guarantee that the following holds for any agent i :

$$\frac{\partial J_i(r_i)}{\partial r_i} \left[\Lambda^* - \mathbb{E}[\lambda_i(s)|s \geq r_i] \right] \geq 0 \geq \lambda_i(r_i) - 2\mathbb{E}[\lambda_i(s)|s \geq r_i] \quad (2.8.29)$$

Then, the inequality given in (2.8.24) is satisfied and the inequality adjusted valuation φ_i is hence weakly increasing.

Part 2: If $J'_i(r_i) \geq 0$, assumption 1 is satisfied.

If $J_i(r_i)$ is weakly monotonic, it will cross 0 at most once. Suppose there exists an r_i^* such that $J_i(r_i^*) = 0$. Because J_i is weakly increasing, setting $\hat{r}_i = r_i^*$ satisfies our requirements. If such a point does not exist, set $\hat{r}_i = 1$.

Part 3: Showing that $\gamma_i(r_i)$ is monotone under the stated assumptions.

The inequality adjusted valuation γ_i reads as follows:

$$\gamma_i = \Lambda_i J_i(r_i) + \Pi_i(r_i) = \Lambda_i \left(r_i - \frac{1 - G_i(r_i)}{g_i(r_i)} \right) + \frac{1 - G_i(r_i)}{g_i(r_i)} \mathbb{E}[\lambda_i(s)|s > r_i] \quad (2.8.30)$$

Taking the derivative of γ_i w.r.t r_i yields:

$$\frac{\partial \gamma_i}{\partial r_i} = \Lambda_i \frac{\partial J_i(r_i)}{\partial r_i} + \left[\frac{\partial}{\partial r_i} \left(\frac{1 - G_i(r_i)}{g_i(r_i)} \right) \right] \mathbb{E}[\lambda_i(s) | s \geq r_i] + \left[\frac{1 - G_i(r_i)}{g_i(r_i)} \right] \frac{\partial \mathbb{E}[\lambda_i(s) | s \geq r_i]}{\partial r_i} \quad (2.8.31)$$

Plugging in our previous results yields that:

$$\frac{\partial \gamma_i}{\partial r_i} = \Lambda_i \frac{\partial J_i(r_i)}{\partial r_i} + \left[\frac{\partial}{\partial r_i} \left(\frac{1 - G_i(r_i)}{g_i(r_i)} \right) \right] \mathbb{E}[\lambda_i(s) | s \geq r_i] + \left[\mathbb{E}[\lambda_i(s) | s \geq r_i] - \lambda_i(r_i) \right] \quad (2.8.32)$$

Thus:

$$\frac{\partial \gamma_i}{\partial r_i} \geq 0 \iff \frac{\partial J_i(r_i)}{\partial r_i} \left[\Lambda_i - \mathbb{E}[\lambda_i(s) | s \geq r_i] \right] \geq \lambda_i(r_i) - 2\mathbb{E}[\lambda_i(s) | s \geq r_i] \quad (2.8.33)$$

Thus, the assumption that $\lambda_i(s)$ is weakly decreasing in s , together with the assumption that $J'_i(r_i) \geq 0$, implies that the LHS is strictly positive. The RHS is negative because $\bar{v}_i^M \leq 2\underline{v}_i^M$. This proves that γ_i is also monotone under the stated assumptions.

2.8.6 Proof of Corollary 3

Part 1: When $\lambda_i(r_i) = 1$, $\gamma_i(r_i) = \varphi_i(r_i) = r_i$ for all r_i .

When $\lambda_i(r_i) = 1$ for all r_i , we have $\Lambda_i = \Lambda^* = \mathbb{E}[\lambda_i(s) | s \geq r_i] = 1$, which implies that $\gamma_i(r_i) = \varphi_i(r_i) = r_i$ for all agents i .

Part 2: When (v_i^K, v_i^M) is i.i.d. for all agents i , the utilitarian optimal allocation rule is the ex post efficient allocation rule.

When (v_i^K, v_i^M) is i.i.d. for all agents i , we have $\varphi_i(r) = \varphi_j(r) = \varphi(r)$ for all i, j . By assumption, $\varphi_i(r_i)$ is strictly increasing. This implies that $\varphi(r)$ is a strictly increasing transformation of r and therefore $r_i \geq r_j$ if and only if $\varphi(r_i) \geq \varphi(r_j)$.

Moreover, the inequality adjusted valuations of all agents will be weakly positive for any $r_i \geq 0$. To see this, note that the i.i.d assumption on (v_i^K, v_i^M) implies that $\Lambda^* = \Lambda_i$ holds for all i . Thus, $\gamma_i(0) = \varphi_i(0) = 0$ holds for all agents. Because the inequality adjusted valuations are monotone, they will be strictly positive for any $r_i > 0$.

Thus, the utilitarian optimal allocation rule is the ex-post efficient rule. The goods will be allocated to the agents with the highest willingnesses-to-pay.

2.8.7 Proof of Proposition 4

Consider the inequality adjusted valuation $\varphi_i(r_i)$ at the lowest possible realization \underline{r}_i :

$$\varphi_i(\underline{r}_i) = \Pi_i(\underline{r}_i) + \Lambda^* J_i(\underline{r}_i) \quad (2.8.34)$$

$$= \frac{\int_{\underline{r}_i}^{\bar{r}_i} \lambda(s) dG_i(s)}{g_i(\underline{r}_i)} + \Lambda^* \left(\underline{r}_i - \frac{1 - G_i(\underline{r}_i)}{g_i(\underline{r}_i)} \right) \quad (2.8.35)$$

$$= \Lambda^* \underline{r}_i + \frac{\Lambda_i - \Lambda^*}{g_i(\underline{r}_i)} \quad (2.8.36)$$

We note that for agent i^* , this expression is weakly positive if and only if $\underline{r}_i \geq 0$. Therefore, the only reason to not allocate the good to agent i^* would be that she has a negative valuation for the good. For the other agents, the expression is weakly positive if and only if

$$\Lambda^* \underline{r}_i + \frac{\Lambda_i - \Lambda^*}{g_i(\underline{r}_i)} \geq 0 \quad (2.8.37)$$

which will generally subject them to rationing unless \underline{r}_i is sufficiently large.

2.8.8 Proof of Proposition 5

For easier reference, we restate this proposition now:

Assume that the marginal utility of money of all agents is non stochastic. Then, it holds that:

1. $\frac{\partial Pr(\varphi_i(r_i) < 0)}{\partial \Lambda^*} \geq 0$ holds for all agents $i \neq i^*$. Thus, when Λ^* increases, the probability with which rationing occurs weakly increases.
2. When $\frac{\partial Pr(\varphi_i(r_i) < 0)}{\partial \Lambda_i} < 0$, a decrease of $\Lambda_i \neq \Lambda^*$ will imply an increase of the probability with which rationing will occur. Note that $\frac{\partial Pr(\varphi_i(r_i) < 0)}{\partial \Lambda_i} < 0$ holds true if the virtual valuation is weakly increasing.

Part 1: An increase of Λ^* weakly increases $Pr(\varphi_i(r_i) < 0)$ for all agents, which leads to an increase in the probability with which rationing occurs.

Firstly, we need to show that:

$$\frac{\partial Pr(\varphi_i(r) < 0)}{\partial \Lambda^*} \geq 0 \quad (2.8.38)$$

Consider any agent i and note the following:

$$\frac{\partial \varphi_i(r)}{\partial \Lambda^*} = J_i(r) \quad \forall r \in (\underline{r}_i, \bar{r}_i) \quad (2.8.39)$$

To understand the effect of an increase in Λ^* on the probability with which agent i is rationed, note firstly that the random variable is r_i .

Firstly, consider realizations of r_i where $\varphi_i(r_i) < 0$ a priori. Because $\varphi_i(r_i) < 0$, it must hold that $J_i(r_i) < 0$ at these realizations of r_i . For these realizations, an increase in Λ^* will thus reduce $\varphi_i(r_i)$, keeping this negative for all these realizations of r_i .

Secondly, consider realizations of r_i where $\varphi_i(r_i) \geq 0$ and $J_i(r_i) \geq 0$ holds true. For these realizations of r_i , the increase in Λ^* will imply a weak increase of $\varphi_i(r_i)$, such that $\varphi_i(r_i) \geq 0$ will still hold after the increase in Λ^* .

Thirdly and finally, consider realizations of r_i where $\varphi_i(r_i) \geq 0$ and $J_i(r_i) < 0$ holds true. For these realizations of r_i , the increase in Λ^* will imply a (weak) decrease of φ_i , which can potentially push those into the negative region, even though they were positive ex ante. This working channel has a weakly positive effect on the probability that this agent is rationed.

Thus, the above arguments imply that $Pr(\varphi_i(r_i) < 0) < 0$ is weakly rising in Λ^* .

It remains to argue why the increase in $Pr(\varphi_i < 0)$ implies an increase in the probability with which rationing occurs. The probability that rationing occurs is the probability that at least $N - m + 1$ agents have negative inequality adjusted valuations. Because the agents' inequality adjusted valuations are independent, any increase of $Pr(\varphi_i(r_i) < 0)$ must lead to an increase in the probability that rationing occurs.

Point 2:

It was previously argued that an increase of $Pr(\varphi_i(r) < 0)$ will imply an increase in the probability with which rationing occurs. This implies the first sentence of this point.

It remains to show the following:

$$\frac{\partial J_i(r_i)}{\partial \Lambda_i} \geq 0 \implies \frac{\partial Pr(\varphi_i(r_i) < 0)}{\partial \Lambda_i} < 0 \quad (2.8.40)$$

Because any agent's marginal utility of money is non-stochastic, their inequality adjusted valuations are:

$$\varphi_i(r_i) = \Lambda_i r_i + (\Lambda^* - \Lambda_i) J_i(r_i) \quad (2.8.41)$$

The derivative of φ_i with respect to Λ_i is:

$$\frac{\partial \varphi_i(r_i)}{\partial \Lambda_i} = r_i + (\Lambda^* - \Lambda_i) \frac{\partial J_i(r_i)}{\partial \Lambda_i} + (-1)J_i(r_i) \quad (2.8.42)$$

$$= r_i + (\Lambda^* - \Lambda_i) \frac{\partial J_i(r_i)}{\partial \Lambda_i} + (-1) \left[r_i - \frac{1 - G_i(r_i)}{g_i(r_i)} \right] \quad (2.8.43)$$

$$= (\Lambda^* - \Lambda_i) \frac{\partial J_i(r_i)}{\partial \Lambda_i} + \frac{1 - G_i(r_i)}{g_i(r_i)} \quad (2.8.44)$$

When the virtual valuation is weakly increasing in Λ_i , the inequality adjusted valuation will uniformly increase in Λ_i . A uniform increase in $\varphi_i(r_i)$ as a result of a change in Λ_i implies that $Pr(\varphi_i(r_i) < 0)$ will fall.

2.8.9 Proof of Proposition 6

We work with the following functions:

$$\Gamma_i(r_i) := Pr\{\max_{j \neq i} \{\gamma_j(r_j)\} \leq \gamma_i(r_i)\} \quad (2.8.45)$$

$$\tau_i(r_i) := \begin{cases} r_i - (1/\Gamma_i(r_i)) \left[\int_{r_i}^{r_i} \Gamma_i(s) ds \right] & r_i > 0 \\ 0 & r_i = 0 \end{cases} \quad (2.8.46)$$

Moreover, note that:

$$\tilde{\tau}_i(b_i) = \begin{cases} \tau_i^{-1}(b_i) & b_i \leq \tau_i(\bar{r}_i) \\ \bar{r}_i & b_i > \tau_i(\bar{r}_i) \end{cases} \quad (2.8.47)$$

$$\tilde{\tau}_N(b_N) = \begin{cases} \tau_N^{-1}(b_N) & b_N \leq \tau_N(\tilde{r}_N) \\ \tilde{r}_N & b_N > \tau_N(\tilde{r}_N) \end{cases} \quad (2.8.48)$$

Agent i wins the auction if and only if:

$$\gamma_i(\tilde{\tau}(b_i)) \geq \max_{j \in \{1, 2, \dots, N\}} \{\gamma_j(\tilde{\tau}_j(b_j))\} \quad (2.8.49)$$

Part 1a: Showing that $\tau_i(r_i)$ is strictly increasing on $[0, \bar{r}_i]$ for $i \leq N - 1$ and on $[0, \tilde{r}_N]$ for agent N when $\gamma_i(r_i)$ is strictly increasing and continuous.

Firstly, recall that the inequality adjusted valuations were defined as follows:

$$\gamma_i(r_i) = \Lambda_i r_i + \frac{\int_{r_i}^{\bar{r}_i} (\lambda_i(s) - \Lambda_i) dG_i(s)}{g_i(r_i)} \quad (2.8.50)$$

When $r_i = \underline{r}_i = 0$, $\gamma_i(\underline{r}_i) = 0$. Because the function γ_i is continuous and strictly increasing, the allocation probability $\Gamma_i(r_i)$ will be strictly positive for any $r_i > 0$.

On $r_i \in (0, \bar{r}_i)$, the derivative of the function τ_i w.r.t. r_i is:

$$\frac{\partial \tau_i(r_i)}{\partial r_i} = 1 - \frac{(\Gamma_i(r_i))(\Gamma_i(r_i)) - (\int_{\underline{r}_i}^{r_i} \Gamma_i(s) ds)(\Gamma_i'(r_i))}{(\Gamma_i(r_i))^2} = \frac{(\int_{\underline{r}_i}^{r_i} \Gamma_i(s) ds)(\Gamma_i'(r_i))}{(\Gamma_i(r_i))^2} \quad (2.8.51)$$

This derivative is strictly positive under the stated specifications. For all agents $i \in \{1, \dots, N-1\}$, the function $\gamma_i(r_i)$ will be strictly increasing in r_i , which implies that $\Gamma_i(r_i)$ will also be strictly increasing in r_i .

For agent N , the function $\gamma_N(r_N)$ is also strictly increasing - but this will only strictly increase the allocation probability when $\gamma_N(r_N) < \bar{\gamma}_{N-1}$, because the allocation probability $\Gamma_N(r_N)$ is 1 for any type above this.

Part 1b: For agent N , the function $\tau_N(r_N)$ equals $\tau_N(\tilde{r}_N)$ (i.e. is flat) for any $r_N \geq \tilde{r}_N$.

To see this, note that $\Gamma_N(r_N) = 1$ for any $r_N \geq \tilde{r}_N$. Thus, the function $\tau_N(r_N)$ becomes:

$$\tau_N(r_N) = r_N - \frac{\int_{\underline{r}_N}^{r_N} \Gamma_N(s) ds}{1} = \int_{\underline{r}_N}^{r_N} s \Gamma_N'(s) ds = \int_{\underline{r}_N}^{\tilde{r}_N} s \Gamma_N'(s) ds \quad (2.8.52)$$

Part 1c: The function τ_i is continuous at $r_i = 0$.

To see this, note that this function can be written as follows:

$$\tau_i(r_i) = \frac{\int_0^{r_i} s \Gamma_i'(s) ds}{\Gamma_i(r_i)} \quad (2.8.53)$$

Both terms converge to zero from the top. Applying L'Hopital's rule yields that:

$$\lim_{r_i \downarrow 0} \tau_i(r_i) = \lim_{r_i \downarrow 0} \frac{\int_0^{r_i} s \Gamma_i'(s) ds}{\Gamma_i(r_i)} = \lim_{r_i \downarrow 0} \frac{r_i \Gamma_i'(r_i)}{\Gamma_i'(r_i)} = 0$$

Part 2: Auction equilibrium - assuming that $\bar{\gamma}_{N-1} = \bar{\gamma}_N$.

Under this assumption, the structure of the functions $\tilde{\tau}_i$ is identical for all agents.

We first show that the mechanism we have described induces a social choice function $c(r_i, r_{-i}) = (x(r), t(r))$ that is incentive compatible.

When all agents bid according to the rule $b_i(r_i) = \tau_i(r_i) \leq \tau_i(\bar{r}_i)$, it holds that $\tilde{\tau}_i(\tau_i(r_i)) = r_i$. Then, the interim allocation probabilities induced by this mechanism, namely $X_i(r_i)$, will be equal to $\Gamma_i(r_i)$, which is monotone under our assumptions.

Moreover, the implied transfer rule will satisfy the integrability constraint. To see this, note that any agent that bids b_i makes the expected payment $-b_i\Gamma_i(\tilde{\tau}_i(b_i))$ in equilibrium. When bidding according to $b_i(r_i) = \tau_i(r_i)$, it thus holds that:

$$\tau_i(0) = 0 \implies U_i(\underline{r}_i) = U_i(0) = 0 \quad (2.8.54)$$

Moreover, the interim expected transfers are given by $T_i(r_i) = -\tau_i(r_i)\Gamma_i(r_i)$

The integrability condition is thus satisfied for the implied social choice function because:

$$r_i X_i(r_i) + T_i(r_i) = U_i(\underline{r}_i) + \int_{\underline{r}_i}^{r_i} X_i(s) ds \quad (2.8.55)$$

$$\iff$$

$$r_i \Gamma_i(r_i) - \tau_i(r_i) \Gamma_i(r_i) = \int_{\underline{r}_i}^{r_i} \Gamma_i(s) ds \iff \tau_i(r_i) = r_i - \frac{\int_{\underline{r}_i}^{r_i} \Gamma_i(s) ds}{\Gamma_i(r_i)} \quad (2.8.56)$$

This establishes that the social choice function induced when all agents bid according to $b_i(r_i) = \tau_i(r_i)$ is incentive compatible.

Thus, it is an equilibrium that all agents bid according to this rule. Consider an agent i and suppose that all other agents $-i$ bid according to $b_{-i}(r_{-i}) = \tau_{-i}(r_{-i})$.

By the intermediate value theorem, we have the following: For any bid $b_i \in [0, \tau_i(\bar{r}_i)]$, there exists an $r_i \in [0, \bar{r}_i]$ such that $\tau_i(r_i) = b_i$. Because the social choice function $c(r)$ is incentive compatible, there can be no profitable deviation in the range $[0, \tau_i(\bar{r}_i)]$.

This is because any such bid \hat{b}_i would be associated with an \hat{r}_i such that $\tau_i(\hat{r}_i) = \hat{b}_i$. Thus, this deviation would generate the outcome $c(\hat{r}_i, r_{-i})$ as defined by the social choice function, which cannot make agent i better off.

Now consider possible deviations into the range $b_i > \tau_i(\bar{r}_i)$. The allocation probabil-

ity would be the same as when bidding $b_i = \tau_i(\bar{r}_i)$, but the payment would be higher upon winning the auction - thus, this deviation can also not be strictly profitable.

Part 3: Auction equilibrium - assuming that $\bar{\gamma}_{N-1} < \bar{\gamma}_N$.

Suppose that all agents bid according to $\tau_i(r_i)$. Then, the interim allocation probabilities $X_i(r_i)$ will be equal to $\Gamma_i(r_i)$ for any agents. To see this, consider an agent $i \leq N-1$, for which $X_i(r_i)$ becomes:

$$X_i(r_i) = Pr[\gamma_i(r_i) \geq \max_{j \in \{1,2,\dots,N\}} \{\gamma_j(\tilde{\tau}_i(b_j))\}] \quad (2.8.57)$$

$$= Pr[\gamma_i(r_i) \geq \max_{j \in \{1,2,\dots,N\}} \{\gamma_j(\tilde{\tau}_i(b_j))\} \wedge r_N \geq \tilde{r}_N] + [\gamma_i(r_i) \geq \max_{j \in \{1,2,\dots,N\}} \{\gamma_j(\tilde{\tau}_i(b_j))\} \wedge r_N < \tilde{r}_N] \quad (2.8.58)$$

$$= [\gamma_i(r_i) \geq \max_{j \in \{1,2,\dots,N\}} \{\gamma_j(r_j)\} \wedge r_N < \tilde{r}_N] := \Gamma_i(r_i) \quad (2.8.59)$$

The arguments for the interim allocation probabilities of agent N are analogous. These functions are weakly monotone for all agents.

Moreover, previous results establish that the transfer rule will satisfy integrability. Thus, when all agents bid according to $b_i(r_i) = \tau_i(r_i)$, the resulting social choice function will be incentive compatible.

Consider any agent i . By incentive compatibility, there can be no profitable deviations into the region $[0, \tau_i(\bar{r}_i)]$ and any deviation above this is dominated by deviating to $\tau_i(\bar{r}_i)$.

Now consider agent N . In general, any bid $b_N > \tau_N(\tilde{r}_N)$ is dominated by bidding $b_N = \tau_N(\tilde{r}_N)$. At both these bids, the interim allocation probability is 1, but the interim payment is rising in the b_N .

If $r_N \leq \tilde{r}_N$, incentive compatibility implies that no other bid in $[0, \tau_N(\tilde{r}_N)]$ can yield a better outcome than bidding according to $b_N(r_N) = \tau_N(r_N)$.

If $r_N > \tilde{r}_N$, the best possible bid will be $b_N = \tau_N(\tilde{r}_N)$.

To see this, note the following: For an agent with $r_N = \tilde{r}_N$, the utility of bidding $\tau_N(\tilde{r}_N)$

is:

$$U_N(\tau_N(\tilde{r}_N); \tilde{r}_N) = (\tilde{r}_N - \tau_N(\tilde{r}_N))\Gamma_N(\tau_N(\tilde{r}_N)) \quad (2.8.60)$$

This must be weakly greater than the utility of any other bid in $[0, \tau_N(\tilde{r}_N)]$ by incentive compatibility. For an agent with type $r_N > \tilde{r}_N$, the utility of bidding $\tau_N(\tilde{r}_N)$ is:

$$U_N(\tau_N(\tilde{r}_N); r_N) = (r_N - \tau_N(\tilde{r}_N))\Gamma_N(\tau_N(\tilde{r}_N)) \quad (2.8.61)$$

The utility of making any other bid $\tau_N(\hat{r}_N) < \tau_N(\tilde{r}_N)$ is:

$$U_N(\tau_N(\hat{r}_N); r_N) = (r_N - \tau_N(\hat{r}_N))\Gamma_N(\tau_N(\hat{r}_N)) \quad (2.8.62)$$

Note that $r_N - \tau_N(\hat{r}_N) > 0$ must hold for $r_N > \tilde{r}_N$, since the following chain of inequalities must be satisfied:

$$r_N - \tau_N(\hat{r}_N) > r_N - \tau_N(\tilde{r}_N) > \tilde{r}_N - \tau_N(\tilde{r}_N) \geq 0 \quad (2.8.63)$$

Thus, agent N with type $r_N > \tilde{r}_N$ will optimally bid $b_N = \tau_N(\tilde{r}_N)$ if and only if:

$$U_N(\tau_N(\tilde{r}_N); r_N) \geq U_N(\tau_N(\hat{r}_N); r_N) \iff \frac{r_N - \tau_N(\tilde{r}_N)}{r_N - \tau_N(\hat{r}_N)} \geq \frac{\Gamma_N(\tau_N(\hat{r}_N))}{\Gamma_N(\tau_N(\tilde{r}_N))} \quad (2.8.64)$$

These inequalities must be satisfied for agent N with type $r_N = \tilde{r}_N$. Thus, we can write:

$$\frac{\tilde{r}_N - \tau_N(\tilde{r}_N)}{\tilde{r}_N - \tau_N(\hat{r}_N)} \geq \frac{\Gamma_N(\tau_N(\hat{r}_N))}{\Gamma_N(\tau_N(\tilde{r}_N))} \quad (2.8.65)$$

Consider the following function:

$$LHS(r_N) := \frac{r_N - \tau_N(\tilde{r}_N)}{r_N - \tau_N(\hat{r}_N)} \quad (2.8.66)$$

$$\implies \frac{\partial LHS(r_N)}{\partial r_N} = \frac{(r_N - \tau_N(\hat{r}_N))(1) - (r_N - \tau_N(\tilde{r}_N))(1)}{(r_N - \tau_N(\hat{r}_N))^2} = \frac{\tau_N(\tilde{r}_N) - \tau_N(\hat{r}_N)}{(r_N - \tau_N(\hat{r}_N))^2} > 0 \quad (2.8.67)$$

$$\implies \frac{r_N - \tau_N(\tilde{r}_N)}{r_N - \tau_N(\hat{r}_N)} > \frac{\tilde{r}_N - \tau_N(\tilde{r}_N)}{\tilde{r}_N - \tau_N(\hat{r}_N)} \geq \frac{\Gamma_N(\tau_N(\hat{r}_N))}{\Gamma_N(\tau_N(\tilde{r}_N))} \quad \forall r_N > \tilde{r}_N \quad (2.8.68)$$

This implies that any agent with $r_N > \tilde{r}_N$ would also prefer the bid $\tau_N(\tilde{r}_N)$ over any other bid $\tau_N(\hat{r}_N) < \tau_N(\tilde{r}_N)$. This completes the proof.

2.8.10 Proof of Proposition 7

We work with the following functions:

$$\Phi_i(r_i) := Pr\{\max_{j \neq i} \{\varphi_j(r_j)\} \leq \varphi_i(r_i)\} \quad (2.8.69)$$

$$\beta_i(r_i) = \begin{cases} r_i - \frac{\int_{r_i}^{r_i} \Phi_i(s) ds}{\Phi_i(r_i)} & r_i > r_i^{min} \\ r_i^{min} & r_i \leq r_i^{min} \end{cases} \quad (2.8.70)$$

We can show continuity of this function at r_i^{min} by applying L'Hopital's rule.

We further define the following functions for all agents $i \leq N - 1$:

$$\tilde{\beta}_i(b_i) = \begin{cases} r_i^{min} & b_i \leq \beta_i(r_i^{min}) \\ \beta_i^{-1}(b_i) & b_i \in (\beta_i(r_i^{min}), \beta_i(\tilde{r}_i)) \\ \tilde{r}_i & b_i \geq \beta_i(\tilde{r}_i) \end{cases} \quad (2.8.71)$$

For agent N , this function is:

$$\tilde{\beta}_N(b_N) = \begin{cases} r_N^{min} & b_N \leq \beta_N(r_N^{min}) \\ \beta_N^{-1}(b_N) & b_N \in (\beta_N(r_N^{min}), \beta_N(\tilde{r}_N)) \\ \tilde{r}_N & b_N \geq \beta_N(\tilde{r}_N) \end{cases} \quad (2.8.72)$$

We have defined \tilde{r}_N such that $\varphi_N(\tilde{r}_N) = \varphi_{N-1}$. The allocation rule is:

$$\varphi_i(\tilde{\beta}(b_i)) \geq \max_{j \in \{1, 2, \dots, N\}} \{\varphi_j(\tilde{\beta}(b_j))\} \quad (2.8.73)$$

The implied auction has a Bayes-Nash equilibrium where all agents bid according to the rule $b_i(r_i) = \beta_i(r_i)$.

As before, one can show that the functions $\beta_i(b_i)$ are strictly increasing on $b_i \in (\beta_i(r_i^{min}), \beta_i(\tilde{r}_i))$, which ensures that a well-defined inverse exists on this interval of bids.

Part 2: Auction equilibrium - assuming that $\bar{\varphi}_{N-1} = \bar{\varphi}_N$.

Under this assumption, the structure of the functions $\tilde{\beta}$ is identical for all agents.

We first show that the mechanism we have described induces a social choice function $c(r_i, r_{-i}) = (x(r), t(r))$ that is incentive compatible.

When all agents i bid according to the rule $b_i(r_i) = \beta_i(r_i) \leq \beta_i(\bar{r}_i)$, it holds that $\tilde{\beta}_i(\beta_i(r_i)) = \beta_i^{-1}(\beta_i(r_i)) = r_i$ for any $r_i \in [r_i^{min}, \bar{r}_i]$. For any such agent with $r_i \leq r_i^{min}$, $\tilde{\beta}_i(b_i) = r_i^{min}$, which implies that the agent will never receive the good.

Thus, the induced interim allocation probabilities $X_i(r_i)$ will be equal to $\Phi_i(r_i)$, which is monotone.

Moreover, the implied transfer rule will satisfy the integrability constraint. To see this, note that any agent that bids b_i makes the expected payment $b_i\Phi_i(\tilde{\beta}_i(b_i))$. When bidding according to $b_i(r_i) = \beta_i(r_i)$, it thus holds that:

$$\beta_i(0) = r_i^{min} \implies U_i(r_i) = U_i(0) = 0 \quad (2.8.74)$$

Since the interim expected transfers are given by $T_i(r_i) = -\beta_i(r_i)\Phi_i(r_i)$, the integrability condition is satisfied for the implied social choice function because:

$$r_i X_i(r_i) + T_i(r_i) = U_i(r_i) + \int_{r_i}^{r_i} X_i(s) ds \quad (2.8.75)$$

$$\iff$$

$$r_i \Phi_i(r_i) - \tau_i(r_i) \Phi_i(r_i) = \int_{r_i}^{r_i} \Phi_i(s) ds \iff \tau_i(r_i) = r_i - \frac{\int_{r_i}^{r_i} \Phi_i(s) ds}{\Phi_i(r_i)} \quad (2.8.76)$$

This establishes that the social choice function induced when all agents bid according to $b_i(r_i) = \tau_i(r_i)$ is incentive compatible.

Having established this, we now show that it is an equilibrium when all agents bid according to the rule $\beta_i(r_i)$ defined above.

Consider any agent with $r_i \leq \bar{r}_i$. Any such agent would have no incentives to bid anything else in the interval $b_i \in (\beta_i(r_i^{min}), \beta_i(\bar{r}_i))$, since the existence of a profitable deviation in this region would violate incentive compatibility.

Bidding anything in the interval $b_i \in [0, \beta_i(r_i^{min})]$ yields zero chance of receiving the good and hence 0 utility - thus, this deviation cannot be profitable either. Bidding anything above $\beta_i(\bar{r}_i)$ is dominated by bidding $\beta_i(\bar{r}_i)$, which cannot be profitable by previous arguments.

Part 3: Auction equilibrium - assuming that $\bar{\varphi}_{N-1} < \bar{\varphi}_N$.

First, note that the ex-interim allocation probabilities will also be equal to $\Phi_i(r_i)$ for any agents.

To see this, consider an agent $i \leq N - 1$ and suppose that all agents bid according to $\beta_i(r_i)$. If $r_i \leq r_i^{min}$, the interim allocation probability is 0, as specified by $\Phi_i(r_i)$. Now suppose $r_i \in (r_i^{min}, \bar{r}_i]$, such that $\tilde{\beta}(\beta(r_i)) = r_i$. By the law of total probability, we can write the interim allocation probability as follows:

$$\begin{aligned} X_i(r_i) &= Pr[\varphi_i(r_i) \geq \max_{j \in \{1, 2, \dots, N\}} \{\varphi_j(\tilde{\tau}_i(b_j))\}] = \\ Pr[\varphi_i(r_i) \geq \max_{j \in \{1, 2, \dots, N\}} \{\varphi_j(\tilde{\beta}_i(b_j))\} \wedge r_N \geq \tilde{r}_N] &+ [\varphi_i(r_i) \geq \max_{j \in \{1, 2, \dots, N\}} \{\varphi_j(\tilde{\beta}_i(b_j))\} \wedge r_N < \tilde{r}_N] \\ &= \\ [\varphi_i(r_i) \geq \max_{j \in \{1, 2, \dots, N\}} \{\varphi_j(r_j)\} \wedge r_N < \tilde{r}_N] &:= \Phi_i(r_i) \end{aligned} \quad (2.8.77)$$

Analogous arguments guarantee that $X_N(r_N) = \Phi_N(r_N)$.

This allocation probability is monotone. Moreover, previous results establish that the transfer rule will satisfy integrability. Thus, when all agents bid according to $b_i(r_i) = \beta_i(r_i)$, the resulting social choice function will be incentive compatible.

Consider an agent $i \leq N - 1$. By incentive compatibility, there can be no deviations into the region $(\beta_i(r_i^{min}), \beta_i(\bar{r}_i)]$. Any deviation above this is dominated by deviating to $\beta_i(\bar{r}_i)$. Any deviation to a bid below $\beta_i(r_i^{min})$ will yield 0 utility and thus cannot be profitable either.

Now consider agent N . In general, any bid $b_N > \beta_N(\tilde{r}_N)$ is dominated by bidding $b_N = \beta_N(\tilde{r}_N)$. Similarly, any bid $b_N \leq \beta_N(r_N^{min})$ cannot be a profitable deviation.

If $r_N \leq \tilde{r}_N$, incentive compatibility implies that no other bid in $(\beta_N(r_N^{min}), \beta_N(\tilde{r}_N)]$ can yield a better outcome - thus, it must be optimal for such agents to bid according to $\beta_N(r_N)$.

If $r_N > \tilde{r}_N$, the best possible bid will be $b_N = \tilde{r}_N$.

For an agent with $r_N = \tilde{r}_N$, the utility of bidding $\beta_N(\tilde{r}_N)$ is:

$$U_N(\beta_N(\tilde{r}_N); \tilde{r}_N) = (\tilde{r}_N - \beta_N(\tilde{r}_N)) \Phi_N(\beta_N(\tilde{r}_N)) \quad (2.8.78)$$

This must be weakly greater than the utility of any other bid in $[0, \beta_N(\tilde{r}_N)]$ by incentive compatibility. For an agent with $r_N > \tilde{r}_N$, the utility of bidding $\beta_N(\tilde{r}_N)$ is:

$$U_N(\beta_N(\tilde{r}_N); r_N) = (r_N - \beta_N(\tilde{r}_N)) \Phi_N(\beta_N(\tilde{r}_N)) \quad (2.8.79)$$

The utility of making any other bid $\beta_N(\hat{r}_N) < \beta_N(\tilde{r}_N)$ is:

$$U_N(\beta_N(\hat{r}_N); r_N) = (r_N - \beta_N(\hat{r}_N)) \Phi_N(\beta_N(\hat{r}_N)) \quad (2.8.80)$$

For an agent with $r_N = \tilde{r}_N$, we have:

$$\frac{\tilde{r}_N - \beta_N(\tilde{r}_N)}{\tilde{r}_N - \beta_N(\hat{r}_N)} \geq \frac{\Phi_N(\beta_N(\hat{r}_N))}{\Phi_N(\beta_N(\tilde{r}_N))} \quad (2.8.81)$$

Consider the following function:

$$LHS(r_N) := \frac{r_N - \beta_N(\tilde{r}_N)}{r_N - \beta_N(\hat{r}_N)} \quad (2.8.82)$$

$$\implies \frac{\partial LHS(r_N)}{\partial r_N} = \frac{(r_N - \beta_N(\hat{r}_N))(1) - (r_N - \beta_N(\tilde{r}_N))(1)}{(r_N - \beta_N(\hat{r}_N))^2} = \frac{\beta_N(\tilde{r}_N) - \beta_N(\hat{r}_N)}{(r_N - \beta_N(\hat{r}_N))^2} > 0 \quad (2.8.83)$$

This implies that any agent with $r_N > \tilde{r}_N$ would also prefer the bid $\beta_N(\tilde{r}_N)$ over any other bid $\beta_N(\hat{r}_N) < \beta_N(\tilde{r}_N)$. Thus, we are done.

2.8.11 Calculations of Bidding Subsidies for the Specific Examples

Part 1: No subsidy example:

Consider the simple two-agent case with deterministic marginal utilities of money. Assume $v_i^K \sim U[0, 1]$ for both agents and suppose that $\Lambda_1 > \Lambda_2$. The inequality adjusted valuations become $\gamma_1 = \Lambda_1 r_1$ and $\gamma_2 = \Lambda_2 r_2$. Thus, we have $\bar{\gamma}_i = 1$ for both agents. Note also that $r_i \sim U[0, 1/\Lambda_i]$ holds for both agents.

To define the bidding subsidies, we first need to compute $\Gamma_i(r_i) = Pr(\gamma_i(r_i) \geq \gamma_j(r_j))$.

Noting that $Pr(r_i \leq x) = \Lambda_i x$ holds for both agents, we can write that:

$$\Gamma_i(s) = Pr(\Lambda_i s \geq \Lambda_j r_j) = Pr(r_j \leq (\Lambda_i/\Lambda_j)s) = \Lambda_i s \quad (2.8.84)$$

Thus, we have that:

$$\tau_i(r_i) = r_i - \frac{\int_0^{r_i} \Lambda_i s ds}{\Lambda_i r_i} = r_i - \frac{[0.5s^2]_0^{r_i}}{r_i} = r_i - \frac{0.5(r_i)^2}{r_i} = 0.5r_i \quad (2.8.85)$$

Moreover, note further that:

$$\tau_i^{-1}(y) = 2y \quad (2.8.86)$$

Agent i receives the good under our allocation rule if and only if:

$$\gamma_i(\tau_i^{-1}(b_i)) \geq \gamma_j(\tau_j^{-1}(b_j)) \iff \Lambda_i(2b_i) \geq \Lambda_j(2b_j) \iff \frac{\Lambda_i}{\Lambda_j} b_i \geq b_j \quad (2.8.87)$$

Part 2: Ex ante budget balance example:

Agent 1 has a valuation $v^K \sim U[0, 1]$ and deterministic utility of money $\Lambda_1 = 1$. Thus, $r_1 \sim U[0, 1]$, $\varphi_1(r_1) = 3r_1 - 1$ and $\varphi_1(r_1) \sim [-1, 2]$. Agent 2 has a valuation $v^K \sim U[0, 2]$ and deterministic utility of money $\Lambda_2 = 2$. Thus $r_2 \sim U[0, 1]$, $\varphi_2(r_2) = 2r_2$ and $\varphi_2 \sim U[0, 2]$. Then, we calculate $\Phi_i(r_i) = Pr\{\max_{j \neq i} \{\varphi_j(r_j)\} \leq \varphi_i(r_i) | r_i\}$. We derive $\Phi_1(r_1) = 0.5(3r_1 - 1)$ and $\Phi_2(r_2) = \frac{2r_2+1}{3}$. Then plugging things in, we get:

$$\beta_1(r_1) = r_1 - \frac{\int_{1/3}^{r_1} 0.5(3s - 1) ds}{0.5(3r_1 - 1)} = 0.5r_1 + \frac{1}{6} \quad (2.8.88)$$

$$\beta_2(r_2) = r_2 - \frac{\int_0^{r_2} \frac{2s+1}{3} ds}{\frac{2r_2+1}{3}} = \frac{r_2^2}{2r_2 + 1} \quad (2.8.89)$$

Taking the inverse of $\beta_1(r_1)$ for $r_1 \in [1/3, 1]$ yields:

$$y = 0.5\beta_1^{-1}(y) + \frac{1}{6} \iff \beta_1^{-1}(y) = 2y - \frac{1}{3} \quad (2.8.90)$$

The inverse of $\beta_2(r_2)$, which is defined for all r_2 , must solve the following:

$$y = \frac{\beta_2^{-1}(y)^2}{2\beta_2^{-1}(y) + 1} \iff (\beta_2^{-1}(y))^2 - 2y\beta_2^{-1}(y) - y = 0 \quad (2.8.91)$$

Applying the quadratic formula yields the following solutions for the inverse:

$$\beta_2^{-1}(y) = \frac{2y + / - \sqrt{4y^2 + 4y}}{2} = \frac{2y + / - 2\sqrt{y^2 + y}}{2} = y + / - \sqrt{y^2 + y} \quad (2.8.92)$$

We know that this solution must be in the positive domain. Thus, this inverse must solve:

$$\beta_2^{-1}(y) = y + \sqrt{y^2 + y} \quad (2.8.93)$$

Thus, our allocation rule for $b_1 > 1/3$ is:

$$\begin{aligned} \varphi_1(\beta_1^{-1}(b_1)) > \varphi_2(\beta_2^{-1}(b_2)) &\iff 3\left(2b_1 - \frac{1}{3}\right) - 1 > 2(b_2 + \sqrt{b_2^2 + b_2}) \\ &\iff \end{aligned} \quad (2.8.94)$$

$$6b_1 - 2 > 2b_2 + 2\sqrt{b_2^2 + b_2} \iff b_1 > (1/3)b_2 + (1/3)\sqrt{b_2^2 + b_2} + 1/3 \quad (2.8.95)$$

2.8.12 Derivation of the Uniform v^K , Pareto v^M Example

Let $v^K \sim U[0, 1]$ and $v^M \sim \text{Pareto}(k, x_{\min})$ where v^K and v^M are independent random variables. For the ratio of two independent random variables, the density is given by:

$$g(r) = \int_{-\infty}^{\infty} |v^M| f^K(rv^M) f^M(v^M) dv^M \quad (2.8.96)$$

We note that $v^M \geq 0$ and adjust the integral boundaries to our setting. Then

$$g(r) = \int_{x_{\min}}^{1/r} v^M k x_{\min}^k v^{M-(k+1)} dv^M \quad (2.8.97)$$

$$= \frac{k}{1-k} r^{k-1} x_{\min}^k - \frac{k}{1-k} x_{\min}^k \quad (2.8.98)$$

In the next step we derive $\lambda(r) = \mathbb{E}[v^M|r]$. For this we start by deriving

$$Pr\{v^M \leq u, v^K/v^M \leq r\} = \int_{x_{\min}}^u \int_0^{v^M r} f^M(v^M) f^K(v^K) dv^K dv^M \quad (2.8.99)$$

$$= \frac{k}{k-1} x_{\min}^k r (x_{\min}^{-k+1} - u^{-k+1}) \quad (2.8.100)$$

for $u \leq 1/r$. We restrict to $u \leq 1/r$ as this constitutes the relevant region in which the joint density of v^M and r is strictly positive. We determine this joint density through differentiating twice and determine

$$f(v^M, r) = k x_{\min}^k v^{M-k} \quad (2.8.101)$$

Based on this, we can check the marginal density of r , namely $g(r)$. This is:

$$g(r) = \int_{x_{min}}^{1/r} f(v^M, r) dv^M = \left[\frac{1}{-k+1} k x_{min}^k v^{M-k+1} \right]_{x_{min}}^{1/r} = \frac{k}{1-k} x_{min}^k [r^{k-1} - (x_{min})^{1-k}] \quad (2.8.102)$$

We then derive the conditional density function

$$f(v^M | r) = \frac{(1-k)x_{min}^k v^{M-k}}{r^{k-1} x_{min}^k - x_{min}} \quad (2.8.103)$$

Then we determine $\lambda(r)$

$$\lambda(r) = \int_{x_{min}}^{1/r} v^M \frac{(1-k)x_{min}^k v^{M-k}}{r^{k-1} x_{min}^k - x_{min}} dv^M \quad (2.8.104)$$

$$= \frac{k-1}{k-2} \frac{x_{min}^{2-k} - r^{k-2}}{x_{min}^{1-k} - r^{k-1}} \quad (2.8.105)$$

Integrating up the marginal density $g(r)$ then yields $G(r)$.

$$G(r) = \int g(r) dr = \frac{k}{1-k} x_{min}^k [(1/k)r^k - (x_{min})^{1-k}r] \quad (2.8.106)$$

We can check the marginal density of v^M , which is:

$$f_{v^M}(y) = \int_0^{1/y} f(y, r) dr = \int_0^{1/y} [k x_{min}^k y^{-k}] dr \quad (2.8.107)$$

$$= \left[k x_{min}^k y^{-k} r \right]_0^{1/y} = k x_{min}^k y^{-k} (y)^{-1} = \frac{k(x_{min})^k}{y^{k+1}} \quad (2.8.108)$$

Based on this, we can compute Λ_i , namely:

$$\Lambda = \int_{x_{min}}^{\infty} y f_{v^M}(y) dy = \int_{x_{min}}^{\infty} k(x_{min})^k (y)^{-k} dy = \left[k(x_{min})^k \frac{1}{-k+1} (y)^{-k+1} \right]_{x_{min}}^{\infty} \quad (2.8.109)$$

$$= -k(x_{min})^k \frac{1}{-k+1} (x_{min})^{-k+1} = \frac{k}{k-1} x_{min} \quad (2.8.110)$$

This would not have been necessary, given our knowledge of v^M that we specify a priori, but serves as a nice check of our previous calculations.

One can show that all inequality adjusted valuations satisfy the monotonicity assumptions in this example:




Figure 2.2: Utilitarian optimal allocation rule vs. ex post efficient allocation rule

Chapter 3

Revenue Maximization with Partially Verifiable Information¹

Abstract: I consider a seller selling a good to bidders with two-dimensional private information: their valuation for a good and their characteristic. While valuations are non-verifiable, characteristics are partially verifiable and convey information about the distribution of a bidder's valuation. I derive the revenue-maximizing mechanism and show that it can be implemented by introducing a communication stage before an auction. I show that granting bidders a right to remain anonymous, i.e., to refuse participation in the communication stage, leaves the optimal mechanism unchanged and provides no benefits for the bidders.

Keywords: Mechanism Design, Auctions, Partially Verifiable Types, Communication

JEL Classification: D44, D82, D83

3.1 Introduction

Suppose a seller wants to sell a good to a number of bidders. In his seminal paper, Myerson (1981) showed that the optimal auction employs reserve prices and bidding subsidies. They explicitly depend on the distributions of the valuations of the bidders, which are given exogenously. If the bidders anonymously participate in the auction, there is no directly observable information to condition on, and thus no way to derive distinct distributions of valuations. In this case, the optimal auction treats all bidders equally. Suppose there is some additional, private information that correlates with the bidders' valuations, and that it is possible to (partially) verify this information once it has been volunteered

¹I thank Vitali Gretschko, Volker Nocke and Thomas Troger for their valuable insights. Further, I thank seminar audiences at Mannheim and Toulouse for their feedback. This work was supported by the University of Mannheim's Graduate School of Economic and Social Sciences. Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (Project B01) is gratefully acknowledged.

by bidders. Can the seller benefit from eliciting this additional information from the bidders, even if the information is not directly part of their utility function? Consider some examples of such situations:

Procurement: Consider an auction for the procurement of a good. Typically, the bidders submit their offers to produce the good, and the best offer wins. The bidding strategy of the bidders will generally depend on their cost for the production of the good, which is private information. Suppose the good can be produced with modern machines, at a lower (marginal) cost or with old machines at a higher cost. Can the seller incentivize the bidders to show her their machines? Can she use this information to receive a better offer in the auction?²

Energy Auctions: In energy markets, there are frequent energy balancing auctions to balance out energy supply and demand. Suppose there is an auction in which the bidders offer to supply additional energy. The energy can be produced using gas, coal, solar power or wind. Can the seller incentivize the bidders to offer detailed information about their mode of production?

Wealth in Auctions: Consider an auction for a piece of art. Suppose richer bidders', through more disposable income, are, on average, willing to pay more money for the piece. Additionally, the neighborhood in which they live is a good indicator of wealth. Can the seller elicit the bidders' addresses? Can she use the address information in the auction to generate higher revenue?

All these examples have in common the existence of information which is informative about the valuation of the bidder and is thus relevant to the seller. However, note that this information is typically unobservable to the seller. In procurement, observing the machines that a supplier will use to produce the good is not possible. It is impossible to directly observe what kind of energy source is used to create electricity in energy auctions. It is also not possible to directly observe a bidder's wealth in an art auction. But if a bidder volunteers this information, it may be possible to (partially) verify it. In a procurement auction, the seller cannot verify the machines that are used in production ex ante, but a bidder may invite the seller to show her the machines used to produce the goods. Similarly, an energy provider can provide a detailed production overview of how exactly the energy is produced, that is not publicly available. In the art auction, assume that it is not possible for the seller to observe a bidder's address ex ante. But when she is provided an address by a bidder, she can confirm whether the given address is true or false. For example, the seller could ask the bidder to show her a valid ID document to verify the address. If the bidder lives in a particularly wealthy neighborhood, it is less

²In practice, the buyer often issues a Request for Quote (RfQ) or a Request for Proposal (RfP) to the suppliers. The responses to such requests typically provide such additional information. My paper shows how this information can be used in a procurement auction.

likely that he is poor. If a bidder lives in a comparatively poor neighborhood, it is less likely that he is rich.

Intuitively, this additional information is useful for the seller and can be used to discriminate between the bidders in an optimal mechanism. My paper contributes to two strains of the mechanism design literature: First, the literature of selling a good to buyers in the presence of non-verifiable valuations, as considered in Myerson (1981). Second, the literature of mechanism design with partially verifiable types, as first considered by Green and Laffont (1986).

For the non-verifiable component of my model, I assume that the bidders' preferences are described by a quasi-linear utility function $u_i = \theta_i x_i + t_i$ where θ_i is the valuation of each bidder for the good, x_i denotes the probability with which each bidder receives the good, and t_i are the transfers each bidder receives or pays in the mechanism. I assume that the valuation θ_i is private, non-verifiable information.

For the partially verifiable component of the model, I assume that every bidder has a characteristic $c_i \in C$. Characteristics have no direct impact on a bidder's utility function, but they are informative about the distribution of a bidder's valuation. There is no one-to-one relationship between a certain characteristic and any given valuation. Instead, there is some statistical relationship between characteristics and valuations such that characteristics are informative about the valuations. The characteristic c_i is private, partially verifiable information and C is a finite set containing all possible characteristics. To provide tractable results, I assume that, conditional on the characteristics, it is possible to order the distributions of the valuations according to the *hazard rate order*. Partial verifiability is in the sense of Green and Laffont (1986). For a bidder with characteristic c , there exists a partition of the set C into two sets: first, a set containing the characteristics that cannot be verified to be different from c and second, a set of those characteristics that are verifiably different from c .

After restricting the search for an optimal mechanism to direct mechanisms through an adjusted revelation principle, I show that incentive compatibility boils down to four conditions. The first two are the well-known monotonicity and integrability condition that follow from the application of Milgrom and Segal (2002). The third and fourth conditions relate to the bidders' characteristics. The third condition concerns the ex-interim allocation probability of a bidder that truthfully reports his characteristic. The allocation probability for this first bidder cannot be lower than that of another, second bidder, if the first bidder can present the evidence requested from the second bidder. This condition intuitively follows, as the characteristics of a bidder are not directly relevant for the utility. If this condition is violated, the former bidder can mimic the latter, which would be a profitable deviation. Therefore, bidders can only be treated unequally if one bidder

is asked to present some particular evidence that the other bidder is not able to produce. The fourth condition simply states that a bidder can truthfully report his characteristic and valuation without being asked to present evidence that the bidder is unable to come up with.

Using the integrability condition, the expected revenue generated from the bidders corresponds to their virtual valuation, conditional on their characteristic. However, incentive compatibility now demands the grouping of bidders according to their characteristics. Therefore, the optimal mechanism has to find the optimal grouping structure. Given that the CDFs associated with the characteristics can be ordered according to the hazard rate order, the optimal mechanism groups bidders with characteristics close to one another in terms of the hazard rate order. The exact grouping, however, depends on the exact verifiability structure of the characteristics.

To argue how my mechanism can be integrated into existing auction formats, I show that a two-stage communication plus auction mechanism implements the revenue-maximizing social choice function. In a first stage, the bidders communicate with the seller about their characteristics. The seller then explicitly conditions the auction rules of the auction in the second stage on this communication. I show that the auction in the second stage maximizes the expected revenue, given the equilibrium of the communication stage. Therefore, it is not necessary for the seller to have commitment power regarding the auction rules as a result of the communication stage. This is very desirable for practical applications, as it rules out the temptation to change the rules during the process. Further, in practice, the pre-auction communication can be implemented as easily as asking the bidders to fill out a questionnaire about their characteristics before the auction.

In the baseline model of the two stage implementation, it is not possible for bidders to refuse communication or engage in babbling with the seller. To alleviate concerns about this restriction, I discuss an extension in which I introduce a right to remain anonymous for the bidders. Every bidder can refuse to communicate in the pre-auction communication stage. I show that no bidder benefits from such a right to remain anonymous. I provide an intuitive unraveling result when bidders are granted a right to remain anonymous. Bidders with particularly desirable characteristics intentionally choose to communicate to separate themselves from bidders with less desirable characteristics. This incentive to engage in communication causes an unraveling effect such that in equilibrium, only those bidders with the least desirable characteristics are indifferent between actually remaining anonymous and communicating about their characteristics.

Partially verifiable information presents some technical challenges. As pointed out in

Green and Laffont (1986), the revelation principle does not generally apply to environments with partially verifiable private information. They show that truthful implementation using the revelation principle is only without loss of generality if the structure of the partition of the set of characteristics C satisfies a *nested range condition*.³ If this condition is violated, there are social choice functions that are implementable in a direct mechanism but not truthfully implementable. Singh and Wittman (2001) argue that the nested range condition in Green and Laffont (1986) is too restrictive and excludes many interesting economic applications. In my model, the nested range condition is not necessary and will generally be violated.

To restore the revelation principle for my framework, I follow a more recent approach developed by Strausz (2016). In a methodological contribution, he argues that the failure of the revelation principle in frameworks with partially verifiable types is caused by the modelling approach of Green and Laffont (1986). Then, Strausz (2016) shows how to restore the revelation principle through what he refers to as the *extended environment*: Typically, the social choice function is defined as a mapping from the set of private information into the set of outcomes. In his new approach, he extends social choice functions to also map into the set of partially verifiable characteristics. This addition to the social choice function can be understood as a requirement for the bidders to present *evidence* within the mechanism. Evidence has also been considered by other authors: Kartik and Tercieux (2012) study implementation when bidders can generate evidence for their types at non-prohibitive costs. Ben-Porath and Lipman (2012) extend social choice functions to not just depend on the bidders' preferences, but allow them to submit evidence to support their claims. What sets these papers apart from mine is the general research question: while they consider the general question of implementability, I use their results to characterize the set of implementable social choice functions in my environment. Then, I determine the implementable social choice function that maximizes revenue.

There are other papers that focus on deriving revenue-maximizing mechanisms when information is partially verifiable. Ball and Kattwinkel (2019) derive revenue-maximizing mechanisms for a range of applications in a setting where the principal can use a probabilistic test with binary outcomes to verify the bidders' types. Tests with deterministic outcomes correspond to how partial verifiability is modelled in Green and Laffont (1986), as well as my paper. Generally, their framework allows for tests that are not restricted to deterministic outcomes. However, the authentication rate characterization in Ball and Kattwinkel (2019) reduces to the nested range condition if tests are deterministic. As

³They define the nested range condition as follows: Consider three distinct characteristics $c_1, c_2, c_3 \in C$. Let ϕ denote the partition of C , such that $c' \in \phi(c)$ denotes that a bidder of characteristic c can report characteristic c' . Then the nested range condition is satisfied if: $c_2 \in \phi(c_1)$ and $c_3 \in \phi(c_2) \Rightarrow c_3 \in \phi(c_1)$

my model generally violates the nested range condition, it cannot be nested in their approach. Further, in the auction application within their paper, they consider one dimensional, partially verifiable, private information: the bidders' valuations. My model considers two-dimensional private information instead: non-verifiable valuations and partially verifiable characteristics. This two-dimensional approach stems from a practical concern. Partially verifiable valuations in the auction environment demand some test that allows to verify that a bidder is willing to pay exactly some particular amount of money, say \$100, for a good. However, it seems incredibly difficult to verify a bidders' exact valuation for a good. Verifying some informative characteristics fits a wide range of applications, as pointed out in the examples at the beginning of this paper.

In environments without transfers, Ben-Porath et al. (2014) study the optimal mechanism for a principal who allocates objects to bidders, whose valuation is private information but can be verified at a cost. Li (2020) solves for the optimal mechanism in a setting where the principal can inspect a bidder's report at a cost and impose punishments on false reports. Erlanson and Kleiner (2020) study how a principal should optimally choose between implementing a new policy and maintaining the status quo when information relevant for the decision is privately held by bidders, but can be verified at a cost. However, as all of these papers preclude monetary transfers, they cannot be applied to a bidder-seller situation.

The remainder of the paper is organized as follows: Section 2 presents the model and derives the optimal mechanism. Section 3 discusses the two-stage implementation and the right to remain anonymous. Section 4 concludes.

3.2 Model

3.2.1 Description

Consider a seller (she) and $N \geq 1$ bidders (he) with unit demand. The seller owns $K \geq 1$ units of a homogeneous good. She does not gain utility from the consumption of the goods and is purely interested in revenue maximization. A bidder's utility function over a particular allocation and payment is equal to $\theta_i x_i + t_i$, where $\theta_i \in [\underline{\theta}, \bar{\theta}]$ represents the valuation of each bidder for the good, x_i denotes the probability with which each bidder receives a good and t_i denotes the transfer each bidder receives or pays in the mechanism. I assume that the valuation θ_i is private, non-verifiable information. As a novel feature of my model, every bidder also has a privately known, partially verifiable characteristic c_i where $c_i \in C$ and C is a finite set containing all possible characteristics.

A characteristic c_i does not directly impact a bidder's utility. However, it is informative about the distribution of the bidder's valuation. In an application, characteristics are meant to capture real-life characteristics of the bidders that allow drawing statistical conclusions about the bidders' valuations. For example, the type of machines used to produce a product in the procurement example or the wealth of a bidder in the art auction example. Each characteristic $c_i \in C$ is associated with a CDF $F_{c_i} \in \mathcal{F}$ that governs the distribution of the valuations θ_i . \mathcal{F} denotes the set of all CDFs that are associated with a characteristic in C . When c_i and c'_i are distinct characteristics, their associated CDFs F_{c_i} and $F_{c'_i}$ differ from each other on a set of valuations with strictly positive measure. I assume that all CDFs $F \in \mathcal{F}$ are continuously differentiable and admit strictly positive densities $f > 0$.⁴ The virtual valuation $J(\theta_i)$ of a bidder is defined as $J(\theta_i) = \theta_i - \frac{1-F(\theta_i)}{f(\theta_i)}$ as in Myerson (1981). For simplicity, I assume that all distribution functions are associated with non-decreasing virtual valuations. Conditional on the characteristics c_i and c_j of two distinct bidders, the valuations θ_i and θ_j are distributed independently. There is a common initial prior Δ over the set of characteristics C and hence also over the set \mathcal{F} . The prior Δ assigns a probability $\delta(c_i)$ to each characteristic $c_i \in C$.

For any distribution $F \in \mathcal{F}$, the hazard rate is defined as $\frac{f(\theta)}{1-F(\theta)}$. The hazard rate order \succsim_{hr} relates distributions F and G (noted as $F \succsim_{hr} G$) if $\frac{f(\theta)}{1-F(\theta)} \leq \frac{g(\theta)}{1-G(\theta)}$. I make the following assumption:

Assumption 1 *The hazard rate order \succsim_{hr} establishes a linear order over \mathcal{F} . In particular, for any $F, G \in \mathcal{F}$ it holds that $F \succsim_{hr} G$ or $G \succsim_{hr} F$.*

Without loss of generality, I label the characteristics from $1, 2, \dots, |C|$ s.t. $F_i \succsim_{hr} F_j$ iff $i \geq j$. To build some intuition for the hazard rate order, consider the likelihood ratio order \succsim_{lr} . For any two CDFs F and G let $F \succsim_{lr} G$ iff $\frac{f(\theta_i)}{g(\theta_i)}$ is increasing in θ_i . As an example, any CDFs F and G with increasing density f and decreasing density g satisfy the likelihood ratio order. Note that it is a well established result that $F \succsim_{lr} G \Rightarrow F \succsim_{hr} G$. Therefore, the likelihood ratio order is sufficient for the hazard rate order.⁵ The likelihood ratio order can be interpreted as follows in the context of this paper. Let c_i and c'_i be characteristics such that $F_{c_i} \succsim_{lr} F_{c'_i}$. Then, bidders are more likely to be of characteristic c_i compared to c'_i , the higher the valuation that is considered. There are a variety of situations for which this assumption seems reasonable. Reconsider the examples from the introduction: in the procurement example, it seems intuitive that cheaper (marginal) production costs are more likely for a firm using more modern machines. In the energy auction, marginal costs for energy production using wind or solar power are likely lower

⁴Environments with differing type spaces for the bidders can be approximated through distributions with arbitrarily small densities on certain types in $[\underline{\theta}, \bar{\theta}]$

⁵For a detailed treatment of stochastic orders and further necessary and sufficient conditions for the hazard rate and the likelihood ratio order see Shaked and Shanthikumar (2007)

than using fossil fuels. In the art auction example, wealthier bidders are likely willing to pay more through having more disposable income.

The characteristics are private information. It is not possible for the seller to gather information about a bidder's characteristic ex ante. However, once a bidder reports a particular characteristic, he can submit evidence to support his claim and only then can this evidence be verified by the seller. To model this, I use a correspondence $\phi : C \rightarrow 2^C$. It is a primitive of the model that captures the degree to which characteristics are partially verifiable and whether evidence that is presented by a bidder can be rejected as objectively false or not. Every characteristic c_i is assigned a set of characteristics $\phi(c_i) \subseteq C$. For every reported characteristic \hat{c}_i such that $\hat{c}_i \in \phi(c_i)$, the bidder can produce evidence that cannot be rejected as objectively false. For every reported characteristic $\hat{c}_i \notin \phi(c_i)$ the bidder is unable to produce sufficient evidence to support his claim. I assume that the procedure that is used to judge whether evidence is objectively false is commonly known, hence ϕ is common knowledge. Further, there is no uncertainty in its outcome. I assume that neither the generation of evidence, nor the verification procedure, is associated with any costs for neither the seller nor the bidders. This assumption can be justified in situations where these costs are negligible compared to the value of the goods up for auction. For example, the costs of generating reports are negligible in a multi-million dollar procurement auction.

Note that the set $\phi(c_i)$ explicitly depends on the true characteristic c_i of the bidder. Depending on his true characteristic, a bidder may find it more difficult to produce evidence to back up certain claims \hat{c}_i . To illustrate this point, recall an example from the introduction. Consider the procurement auction and bidder that produces using the most modern machines available on the market. Naturally, he will have a harder time coming up with evidence that he is producing using old machines than a bidder who is actually using old machines. In general, if $\phi(c_i) = C$ for all c_i , the characteristics are completely unverifiable. If $\phi(c_i) = \{c_i\}$, the characteristics are perfectly verifiable, and if $\phi(c_i) \subset C$, the characteristics are partially verifiable. To allow for some tractable results, I assume the following structure regarding the partial verifiability of the characteristics:

Assumption 2 *Truthful disclosure is possible. That is, for all $c \in C$ it holds that $c \in \phi(c)$*

Assumption 3 *For each characteristic $c \in C$ there is a lower bound $\underline{\phi}(c)$ and an upper bound $\bar{\phi}(c)$ such that $\phi(c) = \{c' \in C \mid \underline{\phi}(c) \leq c' \leq \bar{\phi}(c)\}$*

Assumption 4 *The bounds are monotone. Let $c < c'$ be two characteristics, then it holds that $\underline{\phi}(c) \leq \underline{\phi}(c')$ and $\bar{\phi}(c) \leq \bar{\phi}(c')$*

Combining these assumptions highlights the idea that there is a meaningful order included in the labels, such that labels that are further away from each other are more distinct. The further away a particular characteristic is from the bidder's true characteristic, the harder it will be to generate credible evidence for that characteristic. However, it is possible to generate evidence for a characteristic further away from the true characteristic, it must also be possible to generate evidence for a characteristic closer to the truth. Thus, disclosing any characteristic is possible as long as they remain in the bounds set by $\phi(c)$.

Remark: Assumptions 2-4 do not generally guarantee that Green & Laffont's *nested range condition* is satisfied. They define the nested range condition as follows: For any three distinct elements $c_1, c_2, c_3 \in C$, if $c_2 \in \phi(c_1)$ and $c_3 \in \phi(c_2)$ then $c_3 \in \phi(c_1)$. Consider the following example $\phi(c_1) = \{c_1, c_2\}$, $\phi(c_2) = \{c_2, c_3\}$, $\phi(c_3) = \{c_3\}$. It is easy to verify that this example satisfies assumptions 2-4, but violates the nested range condition. If we replace $\phi(c_1)$ with $\phi'(c_1) = \{c_1, c_2, c_3\}$, it is easy to verify that the example satisfies assumptions 2-4 and the nested range condition. This highlights that my model allows for more general partial verifiability structures than those given by the nested range condition.

3.2.2 Mechanisms and the Revelation Principle for Partially Verifiable Types

The goods are allocated through a mechanism. An arbitrary mechanism is denoted by $g = (M, V, x, t)$. Its first component is a set of unverifiable cheap talk messages M . Second, there is a set $V \subseteq C$ of partially verifiable messages. Note that the mechanism does not necessarily have to allow all partially verifiable messages to be sent. The seller may benefit from excluding some messages from the mechanism, such that V may generally be strictly smaller than C . The third component is an allocation rule x that maps all possible combinations of messages into allocations of the goods. The fourth component is a transfer rule t that maps all possible combinations of messages into transfers. A *direct mechanism* is a mechanism in which $M = \Theta$ and $V = C$.

To establish the revelation principle for my model, I follow the approach laid out by Strausz (2016). As the approach is relatively recent, I will briefly present the main definition and result. Using the language of Strausz (2016), I will refer to the environment as defined in Green and Laffont (1986) as the *initial environment*. The environment as defined in Strausz (2016) is referred to as the *extended environment*. Loosely speaking, the environments differ through the introduction of evidence.⁶ Denote by \mathcal{X} the set of all feasible physical allocations of the goods, and \mathcal{T} describes the set of all feasible transfers to

⁶A more detailed explanation, including all the technical definitions, can be found in section 3.5.1 of the appendix.

each of the bidders. The definition of a Bayesian incentive-compatible, direct mechanism is applied to the extended environment:

Definition 3 *A Bayesian incentive compatible, direct mechanism in the extended environment is a tuple $\hat{g} = (x, t, \hat{c})$ with an allocation rule $x : \Theta^N \times C^N \rightarrow \mathcal{X}$, a transfer rule $t : \Theta^N \times C^N \rightarrow \mathcal{T}$ and an evidence rule $\hat{c} : \Theta \times C \rightarrow C$ such that*

$$\theta_i X(\theta_i, c_i) + T(\theta_i, c_i) \geq \theta_i X(\theta'_i, c'_i) + T(\theta'_i, c'_i) - P \cdot \mathbb{1}\{\hat{c}(\theta'_i, c'_i) \notin \phi(c_i)\} \quad (3.2.1)$$

for all $(\theta'_i, c'_i) \in \Theta \times C$, where $X(\theta_i, c_i) = \mathbb{E}_{-i}[x_i(\theta_i, c_i, \theta_{-i}, c_{-i})]$ denotes the ex interim allocation probability and $T(\theta_i, c_i) = \mathbb{E}_{-i}[t_i(\theta_i, c_i, \theta_{-i}, c_{-i})]$ denotes the ex interim expected payment.

Where P denotes a punishment if an agent cannot produce the evidence demanded by the mechanism. The revelation principle can be re-established in this extended environment through standard arguments. Given that the revelation principle holds for this extended environment, it is vital to establish a connection between the extended environment and the initial environment.

Proposition 1 (Strausz (2016)) *Consider the initial environment and its extension. If there exists some mechanism g which implements the social choice function $f : \Theta^N \times C^N \rightarrow \mathcal{X} \times \mathcal{T}$ in the initial environment, then there exists a function $\hat{c} : \Theta \times C \rightarrow C$ such that the extended social choice function $\hat{f}(\cdot) = (f(\cdot), \{\hat{c}(\cdot)\}_{i=1}^N)$ is implementable in a Bayesian incentive compatible, direct mechanism in the extended environment.*

Proof. See appendix. ■

This proposition connects Green & Laffont's initial environment with Strausz' extended environment. It establishes that any social choice function that can be implemented by some mechanism in the initial environment can be truthfully implemented by a direct mechanism in the extended environment with a suitable evidence function. Therefore, I can focus the derivation of the optimal mechanism on direct, incentive-compatible mechanisms in the extended environment without loss of generality. Intuitively, the proposition is made possible through the addition of the evidence function \hat{c} . Using this function, it is possible to take the equilibrium disclosure behavior with respect to the partially verifiable characteristics in any mechanism in the initial environment and define it as the required evidence rule for the extended environment.

3.2.3 Incentive Compatibility and Expected Revenue

The previous section has established a revelation principle for this setup. To proceed, I first offer a full formal description of the maximization problem.

$$\max_{\{x(\theta,c),t(\theta,c),\hat{c}(\theta_i,c_i)\}} \mathbb{E} \left[\sum_{i=1}^N -t_i(\theta, c) \right] \quad (3.2.2)$$

$$\text{s.t. (IC) } \theta_i X(\theta_i, c_i) + T(\theta_i, c_i) \geq \theta_i X(\theta'_i, c'_i) + T(\theta'_i, c'_i) - P \cdot \mathbb{1}\{\hat{c}(\theta'_i, c'_i) \notin \phi(c_i)\} \quad (3.2.3)$$

$$\text{(IR) } \theta_i X(\theta_i, c_i) + T(\theta_i, c_i) \geq 0 \quad (3.2.4)$$

The seller wants to maximize her expected revenue from the allocation of the goods. However, she is restricted to Bayesian incentive-compatible, direct mechanisms that respect individual rationality in the extended environment without loss of generality. As a next step, I further characterize the incentive compatibility constraints.

Proposition 2 *A direct mechanism $\hat{g} = (x, t, \hat{c})$ in the extended environment is Bayesian incentive compatible if and only if the following conditions hold:*

1. *Integrability*

$$\hat{U}(\theta_i, c_i) = \hat{U}(\underline{\theta}, c_i) + \int_{\underline{\theta}}^{\theta_i} X(s, c_i) ds$$

2. *Monotonicity, that is $\theta_i > \theta'_i$ implies $X(\theta_i, c_i) \geq X(\theta'_i, c_i)$*

3. *Optimality with respect to c_i , that is $X(\theta_i, c_i) \geq X(\theta_i, c'_i)$ for all (θ_i, c'_i) such that $\hat{c}(\theta_i, c'_i) \in \phi(c_i)$.*

4. *Feasible evidence for truthful disclosure: $\hat{c}(\theta_i, c_i) \in \phi(c_i)$ for all $\theta_i \in [\underline{\theta}, \bar{\theta}]$ and $c_i \in C$*

Proof. See appendix. ■

Incentive compatibility boils down to four conditions – an integrability condition and a monotonicity condition akin to the literature's standard constraints. The third and fourth conditions are novel and relate to the characteristics. The third condition requires that the ex interim allocation probability of the good for a bidder i with valuation and characteristic (θ_i, c_i) may not be lower than that of a bidder with valuation and characteristic (θ_i, c'_i) if the required evidence $\hat{c}(\theta_i, c'_i)$ can also be submitted by a bidder with characteristic c_i . The intuition for this constraint is that the utility function of the bidders is independent of their characteristics. The utility derived from the consumption of the good solely depends on their valuation for the good. Therefore, an incentive-compatible mechanism cannot assign the good more often to a bidder of a certain characteristic c'_i

compared to a bidder of characteristic c_i if the latter bidder can provide the evidence demanded by the former bidder. The fourth condition states that it must be possible for a bidder to disclose his valuation and characteristic truthfully without being asked to submit evidence that the bidder cannot feasibly submit. It restricts the evidence rule for bidders that are telling the truth, such that they may not be asked for evidence which they cannot generate. By definition of the approach of Strausz (2016), a bidder who discloses evidence that is verifiably false faces a severe punishment.⁷ Thus, a failure of this condition will result in non-truthful disclosure.

Using standard arguments that make use of the integrability condition, the expected transfer conditional on a specific characteristic c_i is given by:

$$\mathbb{E}[T(\theta_i, c_i)|c_i] = \int_{\underline{\theta}}^{\bar{\theta}} X(\theta_i, c_i)J(\theta_i, c_i)f(\theta_i|c_i)d\theta_i \quad (3.2.5)$$

where $J(\theta_i, c_i) = \theta_i - \frac{1-F(\theta_i|c_i)}{f(\theta_i|c_i)}$. This expression for the expected revenue is very similar to the usual condition in the literature, with the difference being that it is conditional on a specific characteristic c_i . In particular, the virtual valuation is calculated using the conditional distribution and density functions. Recall that the probability with which a characteristic c_i occurs is denoted by $\delta(c_i)$. Then I employ the law of iterated expectations to determine the unconditional expected transfers of a bidder.

$$\mathbb{E}[T(\theta_i, c_i)] = \sum_{c_i \in C} \left(\delta(c_i) \int_{\underline{\theta}}^{\bar{\theta}} X(\theta_i, c_i)J(\theta_i, c_i)f(\theta_i|c_i)d\theta_i \right) \quad (3.2.6)$$

The expected total revenue of the seller then equals

$$\mathbb{E} \left[\sum_{i=1}^N T(\theta_i, c_i) \right] = \int_{[\underline{\theta}, \bar{\theta}]^N} \left(\sum_{c \in C^N} \left(\sum_{i=1}^N x_i(\theta, c)J(\theta_i, c_i) \right) \delta(c_1)f(\theta_1|c_1) \cdots \delta(c_N)f(\theta_N|c_N) \right) d\theta \quad (3.2.7)$$

Now, in principle, it is possible to engage in point wise maximization. However, condition 3 of proposition 2 has to be respected.⁸ It is formulated in terms of the ex interim allocation probabilities. Therefore, this introduces some interdependence that has to be addressed first.

⁷For my paper, a sufficiently severe punishment is to exclude such bidders from participating in the auction.

⁸Recall condition 3: Optimality with respect to c_i , that is $X(\theta_i, c_i) \geq X(\theta_i, c'_i)$ for all (θ_i, c'_i) such that $\hat{c}(\theta_i, c'_i) \in \phi(c_i)$.

3.2.4 Deriving the Optimal Mechanism

In this section, I derive the revenue-maximizing mechanism. As argued above, point wise maximization is not quite possible yet due to the interdependence introduced by condition 3 of proposition 2. In the following lemma, I show how to extend this condition onto a point wise basis.

Lemma 1 *Let $g^* = (x^*, t^*, \hat{c}^*)$ be a revenue-maximizing, incentive compatible mechanism. Fix a valuation θ_i and characteristics c_i, c'_i . Then $x_i^*(\theta_i, c_i, \theta_{-i}, c_{-i}) = 1$ and $\hat{c}^*(\theta_i, c_i) \in \phi(c'_i) \Rightarrow x_i^*(\theta_i, c'_i, \theta_{-i}, c_{-i}) = 1$ without loss of generality.*

Proof. See appendix. ■

Lemma 1 shows that the seller cannot gain any additional revenue by trying to separate bidders of different characteristics if bidders of both characteristics can submit the required evidence. This can be proven by examining the revenue of separating bidders of these two characteristics and comparing it to the revenue in which these characteristics are not separated. However, as the proof shows, the revenue of separating the bidders of the different characteristics cannot exceed that of treating the bidders equally. To gain some intuition for why this is true, recall that the characteristics of a bidder are not part of his utility function. The only possibility to treat one bidder differently from the other is through the required evidence. However, if both bidders can produce the required evidence, a different treatment is simply not possible.

Using Lemma 1, a point wise approach to finding the optimal mechanism is possible. We can fix some profile of valuations and then design the optimal evidence rule for the mechanism. The optimal evidence rule then depends on the grouping of different characteristics that it achieves, and the optimal grouping will critically depend on assumption 1.⁹ To see why, recall the definition of the virtual valuation $J_i(\theta_i, c_i)$ of a bidder i conditional on his characteristic c_i :

$$J(\theta_i, c_i) = \theta_i - \underbrace{\frac{1 - F(\theta_i|c_i)}{f(\theta_i|c_i)}}_{\text{inverse hazard rate}} \quad (3.2.8)$$

Now consider two characteristics c_i and c_j such that $F(\theta_i|c_i) \succ_{hr} F(\theta_j|c_j)$. Then comparing the virtual valuations we get:

$$F(\theta_i|c_i) \succ_{hr} F(\theta_j|c_j) \Rightarrow \frac{f(\theta_i|c_i)}{1 - F(\theta_i|c_i)} \leq \frac{f(\theta_j|c_j)}{1 - F(\theta_j|c_j)} \Rightarrow J(\theta_i|c_i) \leq J(\theta_j|c_j) \quad (3.2.9)$$

⁹Recall assumption 1: The hazard rate order \succ_{hr} establishes a linear order over \mathcal{F} . In particular, for any $F, G \in \mathcal{F}$ it holds that $F \succ_{hr} G$ or $G \succ_{hr} F$.

Being able to order the CDFs $F(\theta_i|c_i)$ and $F(\theta_j|c_j)$ using the hazard ratio order, allows a uniform ordering over the virtual valuations associated with those CDFs. Recall that Myerson (1981) established that the virtual valuation of a bidder is the maximum revenue that the seller can extract from a bidder through the allocation of the good. Thus, in the benchmark of commonly known characteristics, the seller prefers the allocation of the good to bidders with characteristics that have lower ranks in the hazard rate order, as they have higher virtual valuations. In the setting of partially verifiable characteristics, however, the seller has to elicit the characteristics of the bidders first. It turns out that the intuition from the common knowledge case carries over to the partially verifiable case. The seller will use the evidence rule to get close to what she would want to do were the characteristics commonly known. This intuition is distilled into the following algorithm:

The optimal mechanism: Fix an arbitrary profile of valuations θ . Then proceed as follows:

1. Set $\tilde{C} = C$, $c_i = 1$.
2. Group all bidders with characteristics $c'_i \in \tilde{C}$ such that $c_i \in \phi(c'_i)$ into a group G_{c_i} . Set $\tilde{C} = C \setminus G_{c_i}$, $c_i = c_i + 1$
3. Repeat step 2 until $\tilde{C} = \emptyset$
4. Calculate the expected virtual valuation $J(\theta_i, G_{c_i})$ for each bidder in each group as

$$J(\theta_i, G_{c_i}) = \frac{1}{\sum_{c'_i \in G_{c_i}} \delta(c'_i)} \sum_{c'_i \in G_{c_i}} \delta(c'_i) J(\theta_i, c'_i) \quad (3.2.10)$$

5. Assign the good to bidder i if and only if $J(\theta_i, G_{c_i}) \geq \max_{M:N-1} J(\theta_j, G_{c_j})$ and $J(\theta_i, G_{c_i}) \geq 0$. Where $\max_{M:N-1} J(\theta_j, G_{c_j})$ denotes the M 'th highest virtual valuation of the other $N - 1$ bidders.

Proposition 3 *Assigning the goods according to the algorithm constitutes the revenue-maximizing, incentive-compatible mechanism.*

Proof. See appendix. ■

The proposition claims two properties of the algorithm: Revenue maximization and incentive compatibility. First, I will provide a discussion why the algorithm maximizes revenue. Recall that the hazard rate order allows us to rank the virtual valuation of the bidders conditional on their characteristic such that $J(\theta_i, c_i = 1) \geq J(\theta_i, c_i = 2) \geq \dots \geq J(\theta_i, c_i = |C|)$. Thus, if characteristics were observable, the seller would prioritize assigning the good to bidders with lower characteristics ceteris paribus. Fix a bidder i

and a profile of valuations and characteristics. Consider the most profitable way for the seller to assign the good to any other bidder $j \neq i$. Now define by \bar{c} the characteristic such that the seller prefers assigning the good to bidder i if $c_i \leq \bar{c}$ and prefers assigning the good to bidder j if $\bar{c} \leq c_j$. For now, suppose that \bar{c} is interior, i.e. that $1 < \bar{c} < |C|$.

If characteristics are not observable, but merely partially verifiable, the seller needs to distinguish the cases in which bidder i has characteristic $c_i < \bar{c}$ from those in which he has characteristic c'_i with $\bar{c} < c'_i$. Note that lemma 1 established that it is only optimal to assign the good to bidder i with characteristic c_i and not characteristic c'_i if $\hat{c}(\theta_i, c_i) \notin \phi(c'_i)$. Since the bounds of the evidence that can be generated are monotone by assumption 4, the only way to distinguish the two characteristics is by requiring evidence such that $\hat{c}(\theta_i, c_i) < \underline{\phi}(c'_i)$. The most efficient way to achieve this is to set $\hat{c}(\theta_i, c_i) = \underline{\phi}(c_i)$ as it is done in the algorithm. Since this procedure generally creates groups of characteristics, the allocation of the good to a bidder then generates revenue equal to the expected virtual valuation of a bidder with valuation θ_i in that particular group. Then the optimal allocation rule assigns the goods to the bidders that have the highest expected virtual valuations depending on their group.

Second, I will discuss why the algorithm is incentive compatible. Recall that there are 4 conditions that characterize an incentive compatible mechanism.

Consider the two standard conditions for incentive compatibility: Integrability is satisfied, as it has been used in deriving the virtual valuation. Monotonicity is satisfied, as the virtual valuations are non-decreasing by assumption and a higher virtual valuation leads to a higher probability of being assigned a good.

Now consider the two conditions that relate to the truthful disclosure of the characteristics. Note that the algorithm assigns the bidders to groups, such as to maximize their expected virtual valuations. Recall the third condition of incentive compatibility: Optimality with respect to c_i , that is $X(\theta_i, c_i) \geq X(\theta_i, c'_i)$ for all (θ_i, c'_i) such that $\hat{c}(\theta_i, c'_i) \in \phi(c_i)$. Given that bidders are assigned to groups such as to maximize their expected virtual valuation, and assignment is determined by virtual valuations, it is obvious that this condition is satisfied. A bidder that would deviate to reporting a different characteristic would lower his probability of receiving the good, and therefore deviation is not optimal. The fourth condition is: Feasible evidence for truthful disclosure: $\hat{c}(\theta_i, c_i) \in \phi(c_i)$ for all $\theta_i \in [\underline{\theta}, \bar{\theta}]$ and $c_i \in C$. Clearly, this condition is satisfied by construction.

Note that the proposition generates two interesting cases as corollaries:

Corollary 4 *If characteristics are perfectly verifiable, each group contains exactly one characteristic and the optimal mechanism is the Myerson auction with heterogeneous priors.*

If characteristics are perfectly verifiable, it holds that $\phi(c_i) = \{c_i\}$ for all $c_i \in C$. No bidder can generate evidence for any other characteristic and participation in the auction is individually rational. Therefore, the bidders reveal their true characteristics and the optimal mechanism is equivalent to Myerson's optimal auction. At the other extreme are completely unverifiable characteristics:

Corollary 5 *If characteristics are completely unverifiable, all characteristics are grouped into a single group and the optimal mechanism is a Myerson auction with symmetric priors.*

If characteristics are completely unverifiable, it holds that $\phi(c_i) = C$ for all $c_i \in C$. Since characteristics themselves are not part of the bidder's utility functions, it is impossible to try and discriminate between them in the mechanism. These corollaries link the assumption of symmetric / heterogeneous priors to partially verifiable characteristics and allow a practical interpretation: If there are observable differences between bidders or unobservable, but partially verifiable characteristics, it is possible to discriminate between bidders in an auction to increase revenue. If, on the other hand, there are no observable differences between bidders and any unobservable characteristics are completely unverifiable, it is not feasible to discriminate between bidders in an auction to increase revenue.

3.3 Two-Stage Implementation & A Right to Remain Anonymous

3.3.1 A Two-Stage Implementation

This section discusses the implementation of the optimal mechanism as a two-stage mechanism. It highlights that it is optimal to extend an existing auction format by including pre-auction communication. In a first stage, the bidders and the seller communicate about the bidders' characteristics. In a second stage, the seller sells the goods in an auction, with rules that explicitly depend on the communication of the first stage. Both stages will be designed to implement the revenue maximizing social choice function implied by the optimal mechanism of the preceding section. In practice, such pre-auction communication could be implemented, for example, by asking the bidders to fill in a questionnaire that inquires about their characteristics before the auction. In procurement, such communication is common place: Procurement projects frequently issue a Request for Proposal (RfP) or a Request for Quote (RfQ). These requests describe the procurement project and solicit responses by prospective suppliers. Within those replies, suppliers describe their proposed solutions to the procurement problem in some detail, effectively communicating about their characteristics in the sense of the model.

Formally, the game consists of a *communication stage* in which the bidders communicate about their characteristics by disclosing some characteristic $\hat{c}_i \in \mathcal{C}$. In the second stage - the *auction stage* - the seller designs the auction rules to explicitly depend on the communication stage. Conditional on some equilibrium beliefs of the communication stage, it is optimal for the seller to use a Myerson auction that makes use of these beliefs as priors.¹⁰ However, bidders anticipate this and take it into consideration when deciding on their optimal communication strategy. Note that truthful disclosure of characteristics does not necessarily constitute equilibrium behavior. Therefore, I introduce beliefs that describe the distribution of the valuations θ_i conditional on some reported characteristic \hat{c}_i . For any disclosure \hat{c}_i , denote the associated belief $\hat{F}(\theta_i|\hat{c}_i)$ with \hat{F}_{c_i} .

In a Perfect Bayesian equilibrium, these beliefs have to be derived given the bidders' strategies using Bayes rule whenever possible. As usual, off-equilibrium path beliefs in Perfect Bayesian equilibria may be arbitrary. However, I will restrict off-path beliefs by a concept called *belief monotonicity*. The beliefs satisfy belief monotonicity, if the beliefs attached to particular characteristics can be ordered using the hazard rate order in the same order as the priors. I show that all beliefs on the equilibrium path naturally satisfy belief monotonicity, and extend this property to the off-equilibrium path beliefs by assumption. This restriction of the off-equilibrium path beliefs corresponds to the seller's ability to choose her preferred equilibrium in the mechanism design framework.

Next, I investigate the incentives that the auction stage creates for the communication stage. To derive these incentives, I will present several helpful results.

Lemma 2 *Suppose the auction stage uses Myerson's mechanism. Let F and G be two CDFs such that $F \succsim_{hr} G$. Then any bidder weakly prefers to be assigned distribution G over F .*

Proof. See appendix. ■

Lemma 2 highlights that the hazard rate order is useful for determining the bidder's optimal behavior. It implies weakly dominant strategies for the bidders in the communication stage: disclose the characteristic with the largest hazard rate possible, i.e., the characteristic with the smallest possible label attached to it. In equilibrium, it may happen that bidders of different true characteristics c_i, c'_i pool on the same characteristic \hat{c}_i . Then the distribution of the valuation θ_i conditional on the report \hat{c}_i is a mixing distribution. To deal with that, I present a useful lemma for mixing distributions:

¹⁰In a Myerson auction with K units of a homogeneous good, a bidder receives the good iff his virtual valuation is positive and among the K highest virtual valuations of the N bidders.

Lemma 3 *Let $\alpha \in [0, 1]$ and F and G be such that $F \succ_{hr} G$. Then for $H = \alpha F + (1-\alpha)G$ we have that $F \succ_{hr} H \succ_{hr} G$*

Proof. See appendix. ■

The hazard rate order is preserved under mixing. Note that the result extends to a mixture of more than two distributions. If there are more than two distributions, the hazard rate order will rate the mixing distribution somewhere between the most and least favorable distribution included in the mixture. The exact order depends on the exact probability weights in the mixture.

The beliefs associated with reporting a certain characteristic are of central importance for understanding the bidders' behavior. Therefore, I investigate the belief structure on and off the equilibrium path more closely. To achieve this, I introduce the notion of *belief monotonicity* formally:

Definition 4 (Belief Monotonicity) *Let \hat{F}_{c_i} and $\hat{F}_{c'_i}$ be two beliefs that are associated with the disclosure of any characteristics c_i and c'_i such that $c_i < c'_i$. The beliefs satisfy belief monotonicity if $\hat{F}_{c'_i} \succ_{hr} \hat{F}_{c_i}$*

Beliefs that satisfy belief monotonicity preserve the initial order of the characteristics and their associated distributions under the hazard ratio order. Note that belief monotonicity has to be satisfied for characteristics disclosed on the equilibrium path in any equilibrium.

Lemma 4 *Any beliefs that are on the equilibrium path satisfy belief monotonicity.*

Proof. See appendix. ■

Intuitively, belief monotonicity on the equilibrium path is driven by the weakly dominant strategies of the bidders to disclose the lowest characteristic possible. Bidders with lower characteristics can report lower characteristics by the monotonicity assumption on the bounds of the partially verifiable messages. Given that belief monotonicity must hold on the equilibrium path, I extend the concept to the off-equilibrium path beliefs by assumption.

Assumption 5 *Belief monotonicity holds off the equilibrium path.*

This assumption, together with lemma 4, establishes that the initial order of the characteristics under the hazard rate order, i.e. $F_{|C|} \succ_{hr} F_{|C|-1} \succ_{hr} \dots \succ_{hr} F_1$ must carry over to the beliefs in equilibrium, that is $\hat{F}_{|C|} \succ_{hr} \hat{F}_{|C|-1} \succ_{hr} \dots \succ_{hr} \hat{F}_1$. This order on the equilibrium beliefs is useful to determine the bidder's behavior when disclosing their characteristics. As lemma 2 established, such beliefs imply a (weakly) dominant disclosure

strategy for each bidder: each bidder should disclose the lowest characteristic possible.

Given these preliminary results, it is straightforward to establish the equilibrium of the two stage game in the following proposition:

Proposition 4 *The communication stage of the optimal mechanism induces equilibrium beliefs about characteristics as follows:*

$$\hat{F}_c = \sum_{c' \in D(c)} \delta(c') F_{c'}$$

where $D(c) := \{c' \in C \mid c = \underline{\phi}(c')\}$, $\hat{F}_c = F_{|C|}$ if $D(c) = \emptyset$ and $\delta(c)$ denotes the mass allocated to $c \in C$ under the initial belief Δ . In equilibrium, any bidder discloses the minimum possible characteristic that he can generate evidence for, i.e. $\hat{c}(c) = \underline{\phi}(c)$

The Myerson auction uses virtual valuations $J(\theta_i, \hat{c}_i)$ that are determined by the disclosed valuations θ_i and characteristic \hat{c}_i as follows:

$$J(\theta_i, \hat{c}_i) = \theta_i - \frac{1 - \hat{F}_{\hat{c}_i}(\theta_i)}{\hat{f}_{\hat{c}_i}(\theta_i)}$$

Proof. See appendix. ■

Note that the disclosure behavior in this two-stage mechanism is the same as the evidence rule of the revenue maximizing direct mechanism. Further, the virtual valuations that are used for allocating the good are defined analogously. Therefore, the two-stage mechanism implements the revenue maximizing social choice function.

As the seller uses the communication stage to increase her revenue, there may be a practical concern: Bidders could refuse to communicate with the seller or attempt to engage in babbling. The next section examines this case more closely.

3.3.2 A Right to Remain Anonymous

This section extends the model by giving each bidder a right to remain anonymous in the communication stage. That is, each bidder is free to remain anonymous or to refuse to communicate. In terms of the model, this equals an addition of a characteristic a to the set of all characteristics C such that $a \in \phi(c_i)$ for every bidder of every characteristic $c_i \in C$. The characteristic a is available to all bidders regardless of their true characteristic, and when a bidder discloses this characteristic a , it is interpreted as remaining anonymous.

At first glance, it seems that such a right to remain anonymous should benefit the bidders. In particular, those with characteristics c_i that are undesirable in terms of the hazard rate order. However, I will show that introducing the right to remain anonymous is inconsequential and provides no benefit to the bidders. While it seems appealing for bidders with particularly undesirable characteristics to choose to remain anonymous, it is not optimal for bidders with more desirable characteristics to pool with them in anonymity. Bidders with more desirable characteristics make the strategic choice to take part in communication to separate themselves from the others that find anonymity more desirable.

This leads to an unraveling effect. Whenever bidders of some characteristic find it optimal to leave anonymity, the beliefs over those that remain anonymous have to be updated. Given the updated beliefs, there are now other bidders among the remaining anonymous bidders that find it desirable to leave anonymity. This logic continues onward, such that in the end, only bidders of the least desirable characteristics will remain. If they cannot pool with other bidders on disclosing a specific characteristic, they will find themselves indifferent in between the disclosure of some characteristic, which is disclosed exclusively by bidders of those characteristics and remaining anonymous.

The presence of such an unraveling effect depends explicitly on the beliefs on and off the equilibrium path. A bidder finds it desirable to separate from anonymity if a characteristic with a more beneficial belief is available. However, consider a situation in which all bidders pool on anonymity. Then, the beliefs attached to any other characteristics are off-equilibrium path beliefs. So far, I restricted those beliefs to follow belief monotonicity. But for the discussion of anonymity, I require some further restrictions for off-equilibrium path beliefs. To see why, note that off-equilibrium path beliefs that attach the worst belief to all characteristics, i.e., $\hat{F}_c = F_{|C|}$ for all $c \in C$ also satisfy belief monotonicity. However, given these off-equilibrium path beliefs, pooling in anonymity is an equilibrium for all bidders.

To address this issue, I start with considering the worst on path equilibrium beliefs that can be attached to a characteristic:

Lemma 5 *The worst on path equilibrium belief \hat{F}_{c_i} that can be sustained for the disclosure of any characteristics c_i such that $c_i = \underline{\phi}(c'_i)$ for some $c'_i \in C$ is*

$$\hat{F}_{c_i} = \sum_{\{c'_i | c_i = \underline{\phi}(c'_i)\}} \delta(c'_i) F_{c'_i}$$

Proof. See appendix. ■

Now I extend this belief structure to the off-path beliefs by assumption.

Assumption 6 *The off path beliefs \hat{F}_{c_i} for any characteristics c_i such that $c_i = \underline{\phi}(c'_i)$ for some $c'_i \in C$ are not worse than*

$$\hat{F}_{c_i} = \sum_{\{c'_i | c_i = \underline{\phi}(c'_i)\}} \delta(c'_i) F_{c'_i}$$

In the sense of the hazard ratio order.

As the last step before establishing the proposition, I must consider the ex-ante belief attached to anonymity. Recall the initial beliefs $F_{|C|} \succ_{hr} F_{|C|-1} \succ_{hr} \dots \succ_{hr} F_1$ and the common initial prior Δ that was associated with the set of possible characteristics C . This common initial prior represents the initial belief about a bidder that is anonymous, such that I assign $F_a = \sum_{c_i \in C} \delta(c_i) F_{c_i}$.

Given these preliminaries, I present a proposition that formalizes the intuition of unraveling and establishes that a right to anonymity is not beneficial for the bidders.

Proposition 5 *Suppose that each bidder has the right to remain anonymous. No bidder benefits from the right to remain anonymous.*

Proof. See appendix. ■

The presence of a right to remain anonymous is inconsequential. Even though anonymity may seem appealing at first glance, in equilibrium it is not. If true anonymity were to occur in the equilibrium, it would imply pooling of the bidders on the choice of remaining anonymous. However, such a pooling behavior is not optimal, as it requires bidders with more favorable characteristics to pool with those of less favorable characteristics. But then bidders with more favorable characteristics have an incentive to separate themselves from the rest, which causes unraveling.

3.4 Conclusion

In this paper, I considered a seller who wishes to sell multiple units of a homogeneous good to a group of bidders. Bidders have privately known, unverifiable valuations and privately known, partially verifiable characteristics. I use Strausz (2016)'s methodological contribution to recover the revelation principle for this framework. The structure of the partially verifiable characteristics that I consider is richer than the nested range condition of Green and Laffont (1986). I have shown that the revenue-maximizing mechanism is a

Myerson auction that groups bidders according to their characteristics and the particular verifiability structure. It can be implemented in two stages: First, a communication stage about the bidders' characteristics, according to which beliefs about the distribution of the bidder's valuations are formed. Second, an auction stage in which these beliefs are used to play Myerson's optimal auction mechanism.

Further, the paper highlighted that introducing a right to remain anonymous for the bidders is inconsequential. If the bidders are allowed to refuse participation in the communication stage, the optimal mechanism is unchanged, and in particular, no bidder benefits from the right to remain anonymous. This is due to an unraveling effect: bidders with beneficial characteristics find it optimal to take part in the communication to avoid pooling with bidders that have less desirable characteristics.

3.5 Appendix for Revenue Maximization with Partially Verifiable Information

3.5.1 Explanation of the Initial and Extended Environment

The *initial environment* describes the environment as explained in Green and Laffont (1986). The *extended environment* describes the environment as explained in Strausz (2016).

Outcomes: In the *initial environment*, the set of outcomes is defined by the combination of physical outcomes and transfers: $\mathcal{X} \times \mathcal{T}$. In the *extended environment*, the set of outcomes is defined by $\mathcal{X} \times \mathcal{T} \times C^N$. In addition to defining the physical allocation of the good and the transfers, the outcome in the extended environment specifies a partially verifiable message for each bidder: evidence.

Message Sets: In the *initial environment*, the set of partially verifiable messages that can be sent in the mechanism depends on the true characteristic c of a bidder and is equal to $\phi(c)$. In the *extended environment*, the set of partially verifiable messages that can be sent is independent of the true characteristic of a bidder and is given by some set $V \subseteq C$.¹¹

Utility Functions: In the *initial environment*, the utility function of a bidder over the

¹¹For technical reasons, it is important to extend the set of messages that can be sent by a bidder in this way. By definition, in a Bayesian game, the set of possible actions of a bidder may not depend on his type.

set of outcomes in the initial environment is given by

$$u_i = \theta_i x_i + t_i \quad (3.5.1)$$

In the *extended environment*, the utility function of a bidder over the set of outcomes in the extended environment is given by

$$\hat{u}_i = \begin{cases} \theta_i x_i + t_i & \text{if } \hat{c} \in \phi(c_i) \\ \theta_i x_i + t_i - P & \text{if } \hat{c} \notin \phi(c_i) \end{cases} \quad (3.5.2)$$

where \hat{c} is some evidence in the form of a partially verifiable message and P is a sufficiently large punishment that ensures that a bidder will not try to submit evidence that can be objectively rejected as false.¹² Clearly, there is effectively no difference between a bidder simply not being able to send a specific message, as in Green and Laffont (1986) and not wanting to send a message that is strictly dominated as in Strausz (2016).

Social Choice Functions: In the *initial environment*, a social choice function is a mapping

$$f : \Theta^N \times C^N \rightarrow \mathcal{X} \times \mathcal{T} \quad (3.5.3)$$

In the *extended environment*, a social choice function is a mapping

$$\hat{f} : \Theta^N \times C^N \rightarrow \mathcal{X} \times \mathcal{T} \times C^N \quad (3.5.4)$$

The first two components of the social choice function are mappings from the private information of the bidders into allocations and transfers. The third component is a departure from the usual definition of a social choice function. It is a mapping from the private information of the bidders into the set of verifiable messages. This third component can be understood as an evidence rule. For each pair of valuations and characteristics (θ_i, c_i) , it assigns some partially verifiable evidence $c_i \in C$, that the bidder has to submit.

3.5.2 Proof of Proposition 1

The proof follows the structure given in Strausz (2016), with some slight adaptations to my framework.

Suppose some mechanism with allocation rule x , and transfer rule t implements f in the initial environment. Consider some bidder i . Then for the tuple (θ_i, c_i) , given the equilibrium strategies of the bidders $-i$, some strategy leading to the outcome $f(\theta, c)$ is optimal

¹²In equilibrium, it is sufficient to punish a bidder that submits objectively false evidence by excluding him from the auction.

and in particular some verifiable message $\hat{c}(\theta_i, c_i) \in \phi(c_i)$ that bidder i sends when outcome $f(\theta, c)$ is reached in equilibrium is optimal. Consider the direct mechanism $\hat{g} = (f, \hat{c})$ with $\hat{c}(\theta_i, c_i)$ being exactly the mapping that describes the part of the optimal strategy with regards to the verifiable messages in the equilibrium of the initial environment given mechanism g . Suppose all bidders $-i$ truthfully reveal their valuations and characteristics and follow the evidence rule \hat{c} . Fix some tuple (θ_i, c_i) . Bayesian incentive compatibility holds for any c'_i s.t. $\hat{c}(\theta_i, c'_i) \notin \phi(c_i)$. Moreover, the optimality of the strategy leading to the implementation of $f(\theta, c)$ and sending the verifiable message $\hat{c}(\theta_i, c_i)$ implies that Bayesian incentive compatibility holds for any (θ'_i, c'_i) such that $\hat{c}(\theta'_i, c'_i) \in \phi(c)$. Therefore, we have incentive compatibility

3.5.3 Proof of Proposition 2

To make the proof more legible, I list conditions 1-4 of the proposition once again:

1. Integrability

$$\hat{U}(\theta_i, c_i) = \hat{U}(\underline{\theta}, c_i) + \int_{\underline{\theta}}^{\theta_i} X(s, c_i) ds$$

2. Monotonicity, that is $\theta_i > \theta'_i$ implies $X(\theta_i, c_i) \geq X(\theta'_i, c_i)$
3. Optimality with respect to c_i , that is $X(\theta_i, c_i) \geq X(\theta_i, c'_i)$ for all (θ_i, c'_i) such that $\hat{c}(\theta_i, c'_i) \in \phi(c_i)$.
4. Feasible evidence for truthful disclosure: $\hat{c}(\theta_i, c_i) \in \phi(c_i)$ for all $\theta_i \in [\underline{\theta}, \bar{\theta}]$ and $c_i \in C$

First, I show that 1, 2, 3 and 4 imply incentive compatibility. Consider some bidder with true valuation and characteristic (θ_i, c_i) and some possible deviations θ'_i with $\theta'_i < \theta_i$ and $c'_i \in C$. Note that condition 4 guarantees that incentive compatibility would trivially be satisfied if $\hat{c}(\theta'_i, c'_i) \notin \phi(c_i)$. Therefore, I can restrict to the case such that $\hat{c}(\theta'_i, c'_i) \in \phi(c_i)$.

Now consider

$$X(\theta_i, c_i)\theta_i + T(\theta_i, c_i) - (X(\theta'_i, c'_i)\theta_i + T(\theta'_i, c'_i)) \quad (3.5.5)$$

$$= X(\theta_i, c_i)\theta_i + T(\theta_i, c_i) - (X(\theta'_i, c'_i)\theta_i + T(\theta'_i, c'_i)) + X(\theta'_i, c'_i)\theta'_i + T(\theta'_i, c'_i) - (X(\theta'_i, c'_i)\theta'_i + T(\theta'_i, c'_i)) \quad (3.5.6)$$

$$= U(\theta_i, c_i) - U(\theta'_i, c'_i) - (\theta_i - \theta'_i)X(\theta'_i, c'_i) \quad (3.5.7)$$

$$\stackrel{1}{=} \int_{\underline{\theta}}^{\theta_i} X(s, c_i)ds - \int_{\underline{\theta}}^{\theta'_i} X(s, c'_i)ds - (\theta_i - \theta'_i)X(\theta'_i, c'_i) \quad (3.5.8)$$

$$\stackrel{3,4}{\geq} \int_{\theta'_i}^{\theta_i} X(s, c_i)ds - (\theta_i - \theta'_i)X(\theta'_i, c_i) \quad (3.5.9)$$

$$\stackrel{2}{\geq} 0 \quad (3.5.10)$$

Note that by condition 4 it holds that $\hat{c}(\theta'_i, c_i) \in \phi(c_i)$ as the set of verifiable messages does not depend on θ'_i but only on c_i and it must be possible for a bidder with the true value-characteristic pair (θ'_i, c_i) to report their valuation truthfully without punishment. But then if both $\hat{c}(\theta'_i, c'_i) \in \phi(c_i)$ and $\hat{c}(\theta'_i, c_i) \in \phi(c_i)$, condition 3 implies that $X(\theta'_i, c_i) \geq X(\theta'_i, c'_i)$.

A similar argument can be made for $\theta'_i > \theta_i$. Consider some bidder with true valuation θ_i and true characteristic c_i and some possible deviations θ'_i with $\theta'_i > \theta_i$ and $c'_i \in C$. Note that incentive compatibility would trivially be satisfied if $\hat{c}(\theta'_i, c'_i) \notin \phi(c_i)$. Therefore, I can restrict to the case such that $\hat{c}(\theta'_i, c'_i) \in \phi(c_i)$. Now consider

$$X(\theta_i, c_i)\theta_i + T(\theta_i, c_i) - (X(\theta'_i, c'_i)\theta_i + T(\theta'_i, c'_i)) \quad (3.5.11)$$

$$= X(\theta_i, c_i)\theta_i + T(\theta_i, c_i) - (X(\theta'_i, c'_i)\theta_i + T(\theta'_i, c'_i)) + X(\theta'_i, c'_i)\theta'_i + T(\theta'_i, c'_i) - (X(\theta'_i, c'_i)\theta'_i + T(\theta'_i, c'_i)) \quad (3.5.12)$$

$$= U(\theta_i, c_i) - U(\theta'_i, c'_i) - (\theta_i - \theta'_i)X(\theta'_i, c'_i) \quad (3.5.13)$$

$$\stackrel{1}{=} - \left[\int_{\underline{\theta}}^{\theta'_i} X(s, c'_i)ds - \int_{\underline{\theta}}^{\theta_i} X(s, c_i)ds - (\theta'_i - \theta_i)X(\theta'_i, c'_i) \right] \quad (3.5.14)$$

$$\stackrel{3,4}{\geq} - \left[\int_{\theta_i}^{\theta'_i} X(s, c'_i)ds - (\theta'_i - \theta_i)X(\theta_i, c'_i) \right] \quad (3.5.15)$$

$$\stackrel{2}{\geq} 0 \quad (3.5.16)$$

Second, I show that IC implies 1, 2, 3, and 4. First, condition 4 obviously has to hold to allow each bidder to report his valuation and type truthfully without being subjected to punishment. If condition 4 failed, then either there would be some pair of misreports (θ'_i, c'_i) that allows the bidder to avoid the punishment, or if one considers the fringe case

where there would be no possible report that allows the bidder to avoid punishment, the individual rationality constraints will end up being violated. Now, consider 2. Note that incentive compatibility implies that

$$X(\theta_i, c_i)\theta_i + T(\theta_i, c_i) \geq X(\theta'_i, c_i)\theta_i + T(\theta'_i, c_i) \quad (3.5.17)$$

$$X(\theta'_i, c_i)\theta'_i + T(\theta'_i, c_i) \geq X(\theta_i, c_i)\theta'_i + T(\theta_i, c_i) \quad (3.5.18)$$

Rearrange

$$(\theta_i - \theta'_i)(X(\theta_i, c_i) - X(\theta'_i, c_i)) \geq 0 \quad (3.5.19)$$

Which yields monotonicity.

Next, the integrability condition, 1. follows from the fact that the utility itself is independent of c_i and the application of Milgrom and Segal (2002). To see condition 3, consider the following: Suppose there is a c'_i such that $\hat{c}(\theta_i, c'_i) \in \phi(c_i)$ and a set of valuations $\tilde{\Theta}_i$ with positive measure such that $X(\theta_i, c_i) < X(\theta_i, c'_i)$. Consider a value θ_i such that $\tilde{\Theta}_i \subseteq [\underline{\theta}, \theta_i]$. I have already shown that IC implies 1, such that the expected utility from truthful reporting equals

$$\int_{\underline{\theta}}^{\theta_i} X(s, c_i) ds < \int_{\tilde{\Theta}} X(s, c'_i) ds + \int_{[\underline{\theta}, \theta_i] \setminus \tilde{\Theta}} X(s, c_i) ds \quad (3.5.20)$$

Which shows that truthful reporting of c_i at values $\theta_i \in \tilde{\Theta}$ is not optimal and is a contradiction.

3.5.4 Proof of Lemma 1

Proof by contradiction. Suppose the statement does not hold. Then there exists c_i, c'_i and some θ_i such that $x_i^*(\theta_i, c_i, \theta_{-i}, c_{-i}) = 1$ and $\hat{c}^*(\theta_i, c_i) \in \phi(c'_i)$ but $x_i^*(\theta_i, c'_i, \theta_{-i}, c_{-i}) = 0$. Consider some $\epsilon > 0$, small, and a ball with radius ϵ around θ_{-i} . Note as all the virtual valuations are continuous, for small enough ϵ the virtual valuations of all bidders $-i$ can be approximated as constant, subject to a bounded error that vanishes as $\epsilon \rightarrow 0$. Since by assumption the mechanism is revenue-maximizing, the allocation $x_i^*(\theta_i, c_i, \theta_{-i}, c_{-i}) = 1$ is optimal at that particular profile and, for small enough ϵ , in a neighborhood around the valuations of the other bidders in the profile. However, then in the whole neighborhood it holds that $x_i(\theta_i, c_i, \theta_{-i}, c_{-i}) = 1$ but $x_i(\theta_i, c'_i, \theta_{-i}, c_{-i}) = 0$. However, this implies $X(\theta_i, c_i) > X(\theta_i, c'_i)$ for that neighborhood. To not violate condition 3 of incentive compatibility, there must be another neighborhood with radius $\eta > 0$, small, around

the valuations of the other bidders for some profile $(\theta_i, c_i, \theta'_{-i}, c'_{-i})$ in which bidder i is awarded the good for the pair (θ_i, c'_i) but not for (θ_i, c_i) .

Consider the profile $(\theta_i, c_i, \theta_{-i}, c_{-i})$. Denote by J the maximum revenue that can be achieved by incentive compatible assignment of the good to another bidder, or possibly through keeping the good. Denote by J' the revenue associated with the best, alternative incentive compatible alternate assignment given the profile $(\theta_i, c_i, \theta'_{-i}, c'_{-i})$. Then the assignment rule $x_i(\theta_i, c_i, \theta_{-i}, c_{-i}) = 1$, $x_i(\theta_i, c'_i, \theta_{-i}, c_{-i}) = 0$, $x_i(\theta_i, c_i, \theta'_{-i}, c'_{-i}) = 0$ and $x_i(\theta_i, c'_i, \theta'_{-i}, c'_{-i}) = 1$ yields more revenue than the assignment of the good to bidder i at both (θ_i, c_i) and (θ_i, c'_i) given (θ'_{-i}, c'_{-i}) and assigning to the best alternative yielding J at (θ_{-i}, c_{-i}) if

$$Pr(B_\epsilon(\theta_{-i}), c_{-i})[f(\theta_i|c_i)\delta(c_i)(J(\theta_i, c_i) + J') + f(\theta_i|c'_i)\delta(c'_i)(J + J(\theta_i, c'_i))] \quad (3.5.21)$$

$$\geq Pr(B_\eta(\theta'_{-i}), c'_{-i})[f(\theta_i|c_i)\delta(c_i)(J + J(\theta_i, c_i)) + f(\theta_i|c'_i)\delta(c'_i)(J + J(\theta_i, c'_i))] \quad (3.5.22)$$

Note that for sufficiently small ϵ and η the approximation errors will be small enough to be negligible. Further, through appropriate choice of ϵ and η it is possible to set $Pr(B_\epsilon(\theta_{-i}), c_{-i}) = Pr(B_\eta(\theta'_{-i}), c'_{-i})$. There are some more subtle details to note. First, for the bidders $-i$ the changed allocation using the alternatives J and J' is incentive compatible by assumption. Second, the changed allocation is incentive compatible for bidder i as the interim allocation probability $X(\theta_i, c_i)$ and $X(\theta_i, c'_i)$ remains unchanged through the appropriate choice of ϵ and η . Thus, the inequality implies

$$J' - J \geq 0 \quad (3.5.23)$$

Now consider an alternate assignment rule, which assigns the good to bidder i given the profile $(\theta_i, c_i, \theta_{-i}, c_{-i})$ and $(\theta_i, c'_i, \theta_{-i}, c_{-i})$, but never assigns the good to bidder i under alternative profile $(\theta_i, c_i, \theta'_{-i}, c'_{-i})$ and $(\theta_i, c'_i, \theta'_{-i}, c'_{-i})$. Then the following inequality must hold

$$f(\theta_i|c_i)\delta(c_i)(J(\theta_i, c_i) + J') + f(\theta_i|c'_i)\delta(c'_i)(J + J(\theta_i, c'_i)) \quad (3.5.24)$$

$$\geq f(\theta_i|c_i)\delta(c_i)(J(\theta_i, c_i) + J') + f(\theta_i|c'_i)\delta(c'_i)(J(\theta_i, c'_i) + J') \quad (3.5.25)$$

$$\Rightarrow J' - J \leq 0 \quad (3.5.26)$$

Note that the only way both of these inequalities can be true at the same time, is if they hold with equality. However, this implies that the revenue of an incentive compatible mechanism that sets $x_i(\theta_i, c_i, \theta_{-i}, c_{-i}) = 1$ while setting $x_i(\theta_i, c'_i, \theta_{-i}, c_{-i}) = 0$ is the same as that of a mechanism that sets $x_i(\theta_i, c_i, \theta_{-i}, c_{-i}) = 1 = x_i(\theta_i, c'_i, \theta_{-i}, c_{-i}) = 1$. Therefore, I can restrict mechanisms to follow the assertion in the lemma without loss of generality.

3.5.5 Proof of Proposition 3

Fix some bidder i with valuation θ_i and characteristic c_i . For now, suppose that characteristics are observable. Note that for any valuation $\theta_i \in \Theta$ the virtual valuations of the bidder can be ranked according to the hazard rate order, i.e. it holds that $J(\theta_i, c_i = 1) \geq J(\theta_i, c_i = 2) \geq \dots \geq J(\theta_i, c_i = |C|)$. Note that this virtual valuation is exactly the revenue that the seller can extract through the allocation of the good to the bidder. Now consider the choice of the seller: assign the good to bidder i or assign the good to some other bidder j with some valuation and characteristic. It is clear that the larger the characteristic c_i , the smaller the virtual valuation of bidder i , and thus the seller may favor allocation of the good to bidder j . There will be some characteristic \bar{c} such that if $c_i \leq \bar{c}$, the seller wants to allocate the good to bidder i , and if $c_i \geq \bar{c}$ the seller wants to allocate the good to bidder j . The larger the amount of revenue that the seller can receive through the allocation of the good to bidder j , the smaller the value \bar{c} .

Now consider what changes if characteristics are unobservable, but partially verifiable. The seller has to find an evidence rule $\hat{c}(\cdot)$ such that she can distinguish the characteristics $c_i \leq \bar{c}$ from the characteristics $\bar{c} \leq c_i$ if possible. Given assumption 4, namely that the upper and lower bounds of the characteristics for which a bidder can produce evidence are monotone in the true characteristics, it is clear that the optimal evidence rule asks every bidder to produce evidence for the lowest possible characteristic that they feasibly can. To see why, consider the two possible situations that may arise: let c_i be the largest characteristic such that $c_i \leq \bar{c}$ and let c'_i be the smallest characteristic such that $\bar{c} \leq c'_i$. If the evidence rule $\hat{c}(\cdot)$ can distinguish c_i from c'_i , that is if $\underline{\phi}(c_i) < \underline{\phi}(c'_i)$, then the seller can implement the same allocation as she would if characteristics were observable. If the evidence rule $\hat{c}(\cdot)$ is unable to distinguish c_i from c'_i , i.e. if $\underline{\phi}(c_i) = \underline{\phi}(c'_i)$, the monotone bounds assumption implies that $\phi(c_i) \subset \phi(c'_i)$. However, then lemma 1 implies that if the seller assign the good to the bidder with characteristic c_i she must also assign it to the bidder with characteristic c'_i in a revenue maximizing mechanism.

Having established the optimal evidence rule $\hat{c}(\cdot)$, it is straightforward to calculate the expected virtual valuation of bidders that have been grouped together into a group G_{c_i} by the evidence rule through

$$J(\theta_i, G_{c_i}) = \frac{1}{\sum_{c'_i \in G_{c_i}} \delta(c'_i)} \sum_{c'_i \in G_{c_i}} \delta(c'_i) J(\theta_i, c'_i) \quad (3.5.27)$$

Then a point wise maximization implies that is optimal to assign the goods to the buyers

with the largest expected virtual valuations as defined above.

3.5.6 Proof of Lemma 2

Recall that the utility of every bidder at their lowest type in Myerson's optimal auction is set to 0. Then the integrability condition of the incentive compatibility constraints in Myerson's optimal auction reads as:

$$U(\theta_i) = \int_{\underline{\theta}}^{\theta_i} X(\hat{c}_i, s) ds \quad (3.5.28)$$

Where $X(\hat{c}, s)$ denotes the interim allocation probability of bidder i who reports characteristic \hat{c} and valuation s . Recall that a bidder receives the good if and only if his virtual valuation $J(\theta_i, F)$, defined by

$$J(\theta_i, F) = \theta_i - \frac{1 - F(\theta_i)}{f(\theta_i)} \quad (3.5.29)$$

is larger than 0 and any of the other bidders' virtual valuations. Let F and G be such that $F \succsim_{hr} G$. It is easily verified that $J(\theta_i, F) \leq J(\theta_i, G)$ for every θ_i . This implies that $X(\theta_i, F) \leq X(\theta_i, G)$ for every θ_i . Therefore, if the bidder is given the choice between choosing F or G , it is (weakly) dominant to choose G over F for every valuation θ_i .

3.5.7 Proof of Lemma 3

Note that a common alternative characterization of the hazard ratio order is $F \succsim_{hr} G$ iff $\frac{1-G(\theta)}{1-F(\theta)}$ is decreasing in θ . Let $H = \alpha F + (1 - \alpha)G$. First, show that $F \succsim_{hr} H$, i.e. that $\frac{1-H(\theta)}{1-F(\theta)}$ is decreasing in θ .

$$\frac{\partial}{\partial \theta} \left(\frac{1 - H(\theta)}{1 - F(\theta)} \right) = \frac{-h(\theta)(1 - F(\theta)) + f(\theta)(1 - H(\theta))}{(1 - F(\theta))^2} \quad (3.5.30)$$

$$= \frac{-(\alpha f(\theta) + (1 - \alpha)g(\theta))(1 - F(\theta)) + f(\theta)(1 - (\alpha F(\theta) + (1 - \alpha)G(\theta)))}{(1 - F(\theta))^2} \quad (3.5.31)$$

This is negative if the numerator is negative, i.e., if

$$-(\alpha f(\theta) + (1 - \alpha)g(\theta))(1 - F(\theta)) + f(\theta)(1 - (\alpha F(\theta) + (1 - \alpha)G(\theta))) \leq 0 \quad (3.5.32)$$

$$\iff \alpha(-f(\theta))(1 - F(\theta)) + f(\theta)(1 - F(\theta)) + (1 - \alpha)(-g(\theta)(1 - F(\theta)) + f(\theta)(1 - G(\theta))) \leq 0 \quad (3.5.33)$$

$$\iff (1 - \alpha)(-g(\theta)(1 - F(\theta)) + f(\theta)(1 - G(\theta))) \leq 0 \quad (3.5.34)$$

Note that the last inequality holds since $F \succ_{hr} G$.

Second, show that $H \succ_{hr} G$. Consider

$$\frac{\partial}{\partial \theta} \left(\frac{1 - G(\theta)}{1 - H(\theta)} \right) = \frac{-g(\theta)(1 - H(\theta)) + h(\theta)(1 - G(\theta))}{(1 - H(\theta))^2} \quad (3.5.35)$$

Again, this is negative if the numerator is negative, that is if

$$-g(\theta)(1 - H(\theta)) + h(\theta)(1 - G(\theta)) \leq 0 \quad (3.5.36)$$

$$\iff -g(\theta)(1 - (\alpha F(\theta) + (1 - \alpha)G(\theta))) + (\alpha f(\theta) + (1 - \alpha)g(\theta))(1 - G(\theta)) \leq 0 \quad (3.5.37)$$

$$\iff \alpha(-g(\theta)(1 - F(\theta)) + f(\theta)(1 - G(\theta))) \leq 0 \quad (3.5.38)$$

Where the inequality holds since $F \succ_{hr} G$.

3.5.8 Proof of Lemma 4

By contradiction. Let c_i and c'_i be the two smallest characteristics that are disclosed on the equilibrium path with $c_i < c'_i$ such that there are two beliefs \hat{F}_{c_i} and $\hat{F}_{c'_i}$ and a valuation θ_i , where $\frac{\hat{f}_{c_i}(\theta_i)}{1 - \hat{F}_{c_i}(\theta_i)} < \frac{\hat{f}_{c'_i}(\theta_i)}{1 - \hat{F}_{c'_i}(\theta_i)}$. Since the beliefs are on the equilibrium path, they have to be formed according to Bayes' rule. Fix this value of θ_i and denote the set of all characteristics that disclose their characteristic as c_i by $D(c_i)$ and the set of all characteristics that disclose their characteristic as c'_i as $D(c'_i)$. As both of the beliefs are on the equilibrium path, neither $D(c_i)$ nor $D(c'_i)$ are empty. Note that by the same argument as used in the proof of lemma 2, any bidder with valuation θ_i prefers to disclose characteristic c'_i over characteristic c_i if possible. However, as the beliefs are equilibrium beliefs, it must be impossible for any bidder to do so, that is for all characteristics $a \in D(c_i)$, it

holds that $c'_i \notin M(a)$.

Now consider any bidder of characteristic c'' with $c''_i > c'_i > c_i$. Suppose that $c''_i \in D(c_i)$. By assumption 3 we know that $c'_i \in M(c''_i)$. However, then bidders of characteristic c''_i should disclose characteristic c'_i by lemma 2. Therefore, for any bidder of characteristic $c''_i > c'_i > c_i$ we know that $c''_i \notin D(c_i)$. Now consider any bidder of characteristic c''_i with $c''_i < c_i < c'_i$. Since c_i and c'_i are the two smallest characteristics that violate belief monotonicity, we know that $\hat{F}_{c_i} \succ_{hr} \hat{F}_{c''_i}$. Since truthful disclosure of the characteristic is possible by assumption, i.e., $c''_i \in M(c''_i)$, any bidder of such a characteristic is better off disclosing their characteristic truthfully rather than disclosing characteristic c_i . Therefore, for any bidder of characteristic c''_i with $c''_i < c_i < c'_i$ it holds that $c''_i \notin D(c_i)$. Thus, the only bidder that will possibly disclose characteristic c_i is the bidder that actually has characteristic c_i and since it is disclosed on the equilibrium path we have that $D(c_i) = \{c_i\}$ and therefore $\hat{F}_{c_i} = F_{c_i}$. Note that however, for $\frac{\hat{f}_{c_i}(\theta_i)}{1-\hat{F}_{c_i}(\theta_i)} < \frac{\hat{f}_{c'_i}(\theta_i)}{1-\hat{F}_{c'_i}(\theta_i)}$ to hold true on the equilibrium path, there must exist some characteristics $c''_i < c_i$ such that $c''_i \in D(c'_i)$. However, if $c'_i \in M(c''_i)$ by assumption 3 and 4 it holds that $c'_i \in \phi(c_i)$, which is a contradiction to $c'_i \notin \phi(c_i)$.

3.5.9 Proof of Lemma 5

Denote by $D(c_i)$ the set of all characteristics that disclose c_i in equilibrium. By belief monotonicity and lemma 2, it holds that for any characteristic $c'_i \in C$ we have that $c_i > \underline{\phi}(c'_i) \Rightarrow c'_i \notin D(c_i)$. Further, for any characteristic c' such that $c_i < \underline{\phi}(c'_i)$ it holds that $c'_i \notin D(c_i)$ as it would be impossible to disclose this characteristic. Therefore, the only characteristics c'_i that can disclose c_i in equilibrium are such that $c_i = \underline{\phi}(c'_i)$. Note that they will also do so, that is $c_i = \underline{\phi}(c'_i) \Rightarrow c'_i \in D(c_i)$. This holds as belief monotonicity and lemma 2 again imply that no bidder that can disclose c_i would like to disclose any characteristic larger than c_i . Therefore, it holds that $D(c) = \{c'_i | c_i = \underline{\phi}(c'_i)\}$ and the equilibrium belief must be equal to the one given in the lemma and, in particular, cannot be worse than that.

3.5.10 Proof of Proposition 4

Note that lemma 4 established that belief monotonicity holds in the equilibrium of the optimal mechanism. Further, assumption 5 establishes belief monotonicity for the off-equilibrium path beliefs. Further, lemma 2 has established that given a choice between two beliefs F and G with $G \succ_{hr} F$, it is (weakly) optimal for a bidder to choose belief F .

Given belief monotonicity on and off the equilibrium path, this is equivalent to choosing the lower characteristic. Hence, disclosing larger characteristics cannot be a profitable deviation. Finally, note that it is easy to verify that all the beliefs chosen on the equilibrium path, i.e., those where $D(c) \neq \emptyset$ satisfy belief monotonicity through the application of lemma 3. This follows, as bidders always disclose the lowest possible characteristic, such that for characteristics $c_i < c'_i$ the set of bidders that disclose c_i , i.e., $D(c_i)$ features lower characteristics than the set $D(c'_i)$.

3.5.11 Proof of Proposition 5

Suppose that a right to remain anonymous is introduced to the two-stage mechanism, that is, suppose there is a characteristic a such that $a \in \phi(c_i)$ for all $c_i \in C$. The goal is to show that the equilibrium belief associated with anonymity, i.e., \hat{F}_a , in the optimal mechanism with the right to remain anonymous cannot be better than the worst equilibrium belief on the equilibrium path in the mechanism without the right to remain anonymous. I achieve this in two steps.

First, I show that the equilibrium belief of any characteristic that is disclosed on the equilibrium path must not be worse than \hat{F}_a . Let $c_i \neq a$ be any characteristic that is disclosed on the equilibrium path of the optimal mechanism, including the right to remain anonymous, and let \hat{F}_{c_i} be the associated belief. Note that the belief associated with a in any optimal mechanism must fit somewhere in the hazard rate order. If a is disclosed on the equilibrium path, this holds by lemma 4, and if it is not disclosed on the equilibrium path, it holds by assumption. Now suppose that $\hat{F}_{c_i} \succ_{hr} \hat{F}_a$. By lemma 2, we know that disclosing a is preferred by any bidder over c_i . However, then c_i will not be disclosed on the equilibrium path, a contradiction. This implies that the belief attached to anonymity has to be (weakly) worse than that of any characteristic $c_i \neq a$ if such a characteristic is disclosed on the equilibrium path

Second, I show that all characteristics c_i such that there exists a characteristic c'_i with $c_i = \underline{\phi}(c'_i)$ are disclosed on the equilibrium path. Note that doing this establishes the claim. Those characteristics are the only ones disclosed on the equilibrium path in the optimal mechanism without the right to remain anonymous. They cannot have worse equilibrium beliefs than anonymity, as shown above. Start, by considering the set of those characteristics $\{c_i \in C | \exists c'_i \in C \text{ s.t. } c_i = \underline{\phi}(c'_i)\}$. Note that the structure of the message sets by assumptions 3 and 4 implies that this set is equal to $\{1, 2, \dots, \bar{c}\}$ for some $\bar{c} \in C$. I iterate through this set starting from the lowest characteristic and show that there is no equilibrium such that the characteristic is not disclosed on the equilibrium path. Begin

with $c = 1$. Assume $c = 1$ is not disclosed on the equilibrium path. Now I consider two cases: First, suppose that no characteristic other than anonymity is disclosed on the equilibrium path. Then the equilibrium belief of anonymity is $\hat{F}_a = \sum_{c \in C} \delta(c)F_c$. Note that by assumption, the off-path belief \hat{F}_1 for $c = 1$ is not worse than $\sum_{\{c' | 1 = \phi(c')\}} \delta(c')F(c')$. However, this implies that $\hat{F}_a \succsim \hat{F}_1$, as the characteristics in $\{c' | 1 = \phi(c')\}$ are a subset of more beneficial characteristics than those in C itself.¹³ Therefore, there is a profitable deviation and no equilibrium.

Second, suppose that some other characteristic $c_i \neq 1$ is disclosed on the equilibrium path. By belief monotonicity, the equilibrium belief attached to this characteristic has a worse position in the hazard rate order than the off-path belief attached to $c_i = 1$. As argued above, the equilibrium belief attached to anonymity must not be better than that of any characteristic disclosed on the equilibrium path. Therefore, any bidder of a characteristic c_i such that $1 \in \phi(c_i)$ will deviate to the disclosure of $c_i = 1$. Together, both points cover all the cases, such that there is no equilibrium in which $c_i = 1$ is not disclosed on the equilibrium path.

Having established this, the next characteristic to iterate through is $c_i = 2$. However, given that $c_i = 1$ must be disclosed on the equilibrium path, it is possible to remove those characteristics $c_i \in C$ with $1 \in \phi(c_i)$ from the consideration and follow the same arguments made for $c_i = 1$. It is possible to follow this line of argumentation all the way up to \bar{c} . Therefore, in any equilibrium including the right to remain anonymous all characteristics $\{c_i \in C | \exists c'_i \in C \text{ s.t. } c_i = \phi(c'_i)\}$ are disclosed on the equilibrium path. Note that these are the only characteristics disclosed on the mechanism's equilibrium path without a right to remain anonymous and that the equilibrium beliefs have remained unchanged. However, the equilibrium belief of remaining anonymous is not better than any of the beliefs on the equilibrium path. Thus, no bidder has benefited from the right to remain anonymous.

3.5.12 Proof that Non-Decreasing Virtual Valuations are Preserved Under Mixing

Let F and G be such that the virtual valuations associated with them are non-decreasing. Let $\alpha \in [0, 1]$ and $H = \alpha F + (1 - \alpha)G$. Then show that

$$J_H(\theta) = \theta - \frac{1 - H(\theta)}{h(\theta)} \quad (3.5.39)$$

¹³Technically the sets could be equal if characteristics are completely unverifiable, however, corollary 1 has shown that in this case no information is transmitted in equilibrium. In this sense, all bidders are already anonymous. A right to remain anonymous would not yield any benefit in any case

Is non-decreasing in θ . Note that McAfee and McMillan (1987) established that the virtual valuation is non-decreasing if and only if $1/(1 - H(\theta))$ is convex. Therefore, consider

$$\frac{\partial^2}{\partial \theta^2} \left(\frac{1}{1 - H(\theta)} \right) = 2h(\theta)^2(1 - H(\theta))^{-3} + h'(\theta)(1 - H(\theta))^{-2} \quad (3.5.40)$$

This is positive if

$$2h(\theta)^2 + h'(\theta)(1 - H(\theta)) \geq 0 \quad (3.5.41)$$

$$\iff \alpha^2(2f(\theta)^2 + f'(\theta)(1 - F(\theta))) + (1 - \alpha)^2(2g(\theta) + g'(\theta)(1 - G(\theta))) + 4\alpha(1 - \alpha)g(\theta)f(\theta) \geq 0 \quad (3.5.42)$$

This inequality holds since F , and G have non-decreasing virtual valuations by assumption and $g, f \geq 0$.

Bibliography

- Abadi, Joseph and Markus Brunnermeier (2018) Blockchain economics, Technical report, National Bureau of Economic Research.
- Abreu, Dilip (1983) *Repeated games with discounting: A general theory and an application to oligopoly*: Princeton University.
- Abreu, Dilip, David Pearce, and Ennio Stacchetti (1986) Optimal cartel equilibria with imperfect monitoring, *Journal of Economic Theory*, 39 (1), 251–269.
- Adhami, Saman, Giancarlo Giudici, and Stefano Martinazzi (2018) Why do businesses go crypto? An empirical analysis of initial coin offerings, *Journal of Economics and Business*, 100, 64–75.
- Akbarpour, Mohammad, Piotr Dworzak, and Scott Duke Kominers (2020) Redistributive allocation mechanisms, *Available at SSRN 3609182*.
- Armstrong, Mark (2006) Competition in two-sided markets, *The RAND Journal of Economics*, 37 (3), 668–691.
- Arrunada, Benito and Luis Garicano (2018) Blockchain: The birth of decentralized governance, *Pompeu Fabra University, Economics and Business Working Paper Series*, 1608.
- Atkinson, Anthony Barnes and Joseph E Stiglitz (1976) The design of tax structure: direct versus indirect taxation, *Journal of Public Economics*, 6 (1-2), 55–75.
- Azar, Jos² and Xavier Vives (2021) General Equilibrium Oligopoly and Ownership Structure, *Econometrica*, 89 (3), 999–1048.
- Bakos, Yannis and Hanna Halaburda (2018) The role of cryptographic tokens and icos in fostering platform adoption, *Available at SSRN 3207777*.
- Ball, Ian and Deniz Kattwinkel (2019) Probabilistic verification in mechanism design, in *Proceedings of the 2019 ACM Conference on Economics and Computation*, 389–390.

- Belleflamme, Paul and Martin Peitz (2021) *The Economics of Platforms*: Cambridge University Press.
- Ben-Porath, Elchanan, Eddie Dekel, and Barton L Lipman (2014) Optimal allocation with costly verification, *American Economic Review*, 104 (12), 3779–3813.
- Ben-Porath, Elchanan and Barton L Lipman (2012) Implementation with partial provability, *Journal of Economic Theory*, 147 (5), 1689–1724.
- Besley, Timothy and Stephen Coate (1991) Public provision of private goods and the redistribution of income, *The American Economic Review*, 81 (4), 979–984.
- Biais, Bruno, Christophe Bisiere, Matthieu Bouvard, and Catherine Casamatta (2019) The blockchain folk theorem, *The Review of Financial Studies*, 32 (5), 1662–1715.
- Borgers, Tilman and Peter Norman (2009) A note on budget balance under interim participation constraints: the case of independent types, *Economic Theory*, 39 (3), 477–489, 10.1007/s00199-008-0347-7.
- Brzustowski, Thomas, A Georgiadis, and Balızs Szentes (2021) Smart Contracts and the Coase Conjecture, Technical report, Working Paper.
- Cabral, Luis (2011) Dynamic price competition with network effects, *The Review of Economic Studies*, 78 (1), 83–111.
- Catalini, Christian and Joshua S Gans (2018) Initial coin offerings and the value of crypto tokens, Technical report, National Bureau of Economic Research.
- Catalini, Christian, Ravi Jagadeesan, and Scott Duke Kominers (2020) Markets for crypto tokens, and security under proof of stake, *Available at SSRN 3740654*.
- Chen, Pei-Yu and Lorin M Hitt (2002) Measuring switching costs and the determinants of customer retention in Internet-enabled businesses: A study of the online brokerage industry, *Information Systems Research*, 13 (3), 255–274.
- Chen, Yan, Igor Pereira, and Pankaj C Patel (2021) Decentralized governance of digital platforms, *Journal of Management*, 47 (5), 1305–1337.
- Choi, JP and DS Jeon (2022) Platform design biases in ad-funded two-sided markets, *The RAND Journal of Economics*.
- Condorelli, Daniele (2013) Market and non-market mechanisms for the optimal allocation of scarce resources, *Games and Economic Behavior*, 82, 582–591, 10.1016/j.geb.2013.08.008.

- Cong, Lin William, Ye Li, and Neng Wang (2021) Tokenomics: Dynamic adoption and valuation, *The Review of Financial Studies*, 34 (3), 1105–1155.
- Cres, Herve, Mich Tvede et al. (2020) Corporate self-regulation of imperfect competition, Technical report, School of Economics, University of East Anglia, Norwich, UK.
- Crippen, Dan L. (2000) The Budget and Economic Outlook, *United States Congressional Budget Office*.
- Diamond, Peter A and James A Mirrlees (1971) Optimal taxation and public production I: Production efficiency, *The American Economic Review*, 61 (1), 8–27.
- Donovan, Sarah and David Bradley (2019) Real Wage Trends, 1979 to 2018, *Congressional Research Service*.
- Dworczak, Piotr, Scott Duke Kominers, and Mohammad Akbarpour (2021) Redistribution through markets, *Econometrica*, 88 (7), 10.2139/ssrn.3143887.
- Eisenhuth, Roland (2019) Reference-dependent mechanism design, *Economic Theory Bulletin*, 7 (1), 77–103, 10.1007/s40505-018-0144-9.
- Erlanson, Albin and Andreas Kleiner (2020) Costly verification in collective decisions, *Theoretical Economics*, 15 (3), 923–954.
- Esteban, Joan and Debraj Ray (2006) Inequality, lobbying, and resource allocation, *American Economic Review*, 96 (1), 257–279, 10.1257/000282806776157533.
- Farrell, Joseph and Paul Klemperer (2007) Coordination and lock-in: Competition with switching costs and network effects, *Handbook of Industrial Organization*, 3, 1967–2072.
- Farrell, Joseph and Garth Saloner (1986) Installed base and compatibility: Innovation, product preannouncements, and predation, *The American Economic Review*, 940–955.
- Fritz, Dana (2021) Der Kindergarten: Diese Unterschiede über die verschiedenen Kindergarten-Konzeptionen sollten Sie kennen, https://www.kita.de/wissen/kitergarten/#2_Die_Traegerschaft_eines_Kindergartens(11.01.2022).
- Fudenberg, Drew and Eric Maskin (1990) Nash and perfect equilibria of discounted repeated games, *Journal of Economic Theory*, 51 (1), 194–206.
- Fudenberg, Drew and Jean Tirole (1991) *Game theory*: MIT press.
- Gahvari, Firouz and Enlinson Mattos (2007) Conditional cash transfers, public provision of private goods, and income redistribution, *American Economic Review*, 97 (1), 491–502, 10.1257/aer.97.1.491.

- Gershkov, Alex, Jacob Goeree, Alexey Kushnir, Benny Moldovanu, and Xianwen Shi (2013) On the Equivalence of Bayesian and Dominant Strategy Implementation, *Econometrica*, 81 (1), 197–220, 10.3982/ECTA10592.
- Goldstein, Itay, Deeksha Gupta, and Ruslan Sverchkov (2019) Utility Tokens as a Commitment to Competition, *Available at SSRN 3484627*.
- Green, Jerry R and Jean-Jacques Laffont (1986) Partially verifiable information and mechanism design, *The Review of Economic Studies*, 53 (3), 447–456.
- Hart, Oliver and John Moore (1988) Incomplete contracts and renegotiation, *Econometrica*, 755–785.
- (1999) Foundations of incomplete contracts, *The Review of Economic Studies*, 66 (1), 115–138.
- Hashimoto, Kazuhiko and Hiroki Saitoh (2010) Domain expansion of the pivotal mechanism, *Social Choice and Welfare*, 34 (3), 455–470, 10.1007/s00355-009-0411-7.
- Howell, Sabrina T, Marina Niessner, and David Yermack (2020) Initial coin offerings: Financing growth with cryptocurrency token sales, *The Review of Financial Studies*, 33 (9), 3925–3974.
- Huberman, Gur, Jacob D Leshno, and Ciamac Moallemi (2021) Monopoly without a monopolist: An economic analysis of the bitcoin payment system, *The Review of Economic Studies*, 88 (6), 3011–3040.
- Huesmann, Katharina (2017) Public Assignment of Scarce Resources under Income Effects, *Working paper*.
- Jullien, Bruno and Alessandro Pavan (2019) Information management and pricing in platform markets, *The Review of Economic Studies*, 86 (4), 1666–1703.
- Kang, Mingshi and Charles Z Zheng (2019) Necessity of Auctions for Redistributive Optimality, Technical report, Working Paper.
- Kartik, Navin and Olivier Tercieux (2012) Implementation with evidence, *Theoretical Economics*, 7 (2), 323–355.
- Katz, Michael L and Carl Shapiro (1985) Network externalities, competition, and compatibility, *The American Economic Review*, 75 (3), 424–440.
- Kazumura, Tomoya, Debasis Mishra, and Shigehiro Serizawa (2020) Mechanism design without quasilinearity, *Theoretical Economics*, 15 (2), 511–544, 10.3982/TE2910.

- Kotowski, Maciej H (2020) First-price auctions with budget constraints, *Theoretical Economics*, 15 (1), 199–237, 10.3982/TE2982.
- Lewis, A. (2021) *The Basics of Bitcoins and Blockchains: An Introduction to Cryptocurrencies and the Technology That Powers Them (Cryptography, Derivatives Investments, Futures Trading, Digital Assets, NFT)*: Mango Media.
- LH-Mainz (2018) Vergabeverfahren von Kita-Plätzen, Landeshauptstadt Mainz, <https://fragdenstaat.de/anfrage/vergabeverfahren-von-kita-plaetzen/>(11.01.2022).
- Li, Jiasun and William Mann (2018) Initial coin offering and platform building, *SSRN Electronic Journal*, 1–56.
- Li, Yunan (2020) Mechanism design with costly verification and limited punishments, *Journal of Economic Theory*, 186, 105000.
- Li, Zhuoxin and Ashish Agarwal (2017) Platform integration and demand spillovers in complementary markets: Evidence from Facebook's integration of Instagram, *Management Science*, 63 (10), 3438–3458.
- Magill, Michael, Martine Quinzii, and Jean-Charles Rochet (2015) A theory of the stakeholder corporation, *Econometrica*, 83 (5), 1685–1725.
- Maskin, Eric and John Riley (1984) Optimal auctions with risk averse buyers, *Econometrica: Journal of the Econometric Society*, 1473–1518.
- McAfee, R Preston and John McMillan (1987) Auctions and bidding, *Journal of Economic Literature*, 25 (2), 699–738.
- Milgrom, Paul and Ilya Segal (2002) Envelope theorems for arbitrary choice sets, *Econometrica*, 70 (2), 583–601.
- Myerson, Roger B (1981) Optimal auction design, *Mathematics of Operations Research*, 6 (1), 58–73.
- Pai, Mallesh M and Rakesh Vohra (2014) Optimal auctions with financially constrained buyers, *Journal of Economic Theory*, 150, 383–425, 10.1016/j.jet.2013.09.015.
- Peitz, Martin, Sven Rady, and Piers Trepper (2017) Experimentation in two-sided markets, *Journal of the European Economic Association*, 15 (1), 128–172.
- Rochet, Jean-Charles and Jean Tirole (2003) Platform competition in two-sided markets, *Journal of the European Economic Association*, 1 (4), 990–1029.
- Saitoh, Hiroki and Shigehiro Serizawa (2008) Vickrey allocation rule with income effect, *Economic Theory*, 35 (2), 391–401, 10.1007/s00199-007-0235-6.

- Saleh, Fahad (2021) Blockchain without waste: Proof-of-stake, *The Review of Financial Studies*, 34 (3), 1156–1190.
- Shaked, Moshe and J George Shanthikumar (2007) *Stochastic orders*: Springer Science & Business Media.
- Shapiro, Carl and Hal R Varian (1998) *Information rules: A strategic guide to the network economy*: Harvard Business Press.
- Singh, Nirvikar and Donald Wittman (2001) Implementation with partial verification, *Review of Economic Design*, 6 (1), 63–84.
- Sockin, Michael and Wei Xiong (forthcoming) Decentralization through tokenization, *Journal of Finance*.
- Strausz, Roland (2016) Mechanism design with partially verifiable information.
- Szabo, Nick (1997) Formalizing and securing relationships on public networks, *First Monday*.
- Teh, Tat-How (2022) Platform governance, *American Economic Journal: Microeconomics*, 14 (3), 213–54.
- TFEU (1957) Treaty on the Functioning of the European Union.
- Weitzman, Martin L (1977) Is the price system or rationing more effective in getting a commodity to those who need it most? *The Bell Journal of Economics*, 517–524.

Curriculum Vitae

Marco Reuter

- 2017 - 2023** University of Mannheim
Ph.D. in Economics
- 2014 - 2016** University of Bonn
Master of Science in Economics
- 2011 - 2014** University of Bonn
Bachelor of Science in Economics