THE UNIVERSITY
of EDINBURGH

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

Essays on Competition, Market Structures and Public Goods

Kimon Doulis

Economics PhD

The University of Edinburgh

2016

Abstract

Chapter one focuses on optimal pricing in markets of consumption chains. These are markets in which one good is necessary for access to further consumption goods. I analyse optimal pricing for different market structures, focusing on the case of an integrated monopolist and the case of separate firms being in competition across markets, but not within markets. I then compare the outcomes of different market structures using basic welfare measures. I show that, compared to the first best allocation, the allocation implemented under the integrated monopolist tends to have significantly lower consumer surplus and larger producer surplus. Aggregate welfare is surprisingly not much smaller under the integrated firm when compared to a welfare maximising allocation. In some settings the integrated monopolist even implements a welfare maximising allocation. The paper explains and highlights how these results depend largely on which assumptions are made about the information available to consumers.

The second chapter contributes towards the existing literatures on lobbying and on media bias by combining and extending features of both. It aims to analyse optimal slanting policies of interest or media groups and their effect on the distribution of public opinion and its evolution over time by introducing an intertemporal model of grassroots lobbying or media bias. I also allow for more general results than existing models by making fewer distributive assumptions and by allowing for further incentives of agents. In the chapter I combine demand and supply side models for bias. A main focus lies on how optimal slanting, the distribution of public opinion and its evolution over time depend on competition. The chapter aims to examine in which circumstances competition in the media market or the existence of multiple rival lobby groups can be detrimental. It shows how this can be the case because competition can create an incentive to split the public up and cater only to the own market. This can lead to a loss of the middle ground and increased dispersion of public opinion.

The third chapter aims to extend the existing literature on the (in)efficiencies of voluntary contribution mechanisms for public goods. The existing body of research tries to analyse how group size affects the outcomes of such mechanisms asymptotically, while I also focus on results for given group sizes and the effect of the level of group heterogeneity in combination with group size. Agents are ex post heterogeneous in the existing literature; I also allow for them to be heterogeneous ex ante. This means that agents do not only have different valuations for the public good ex post, but different agents are also perceived differently by other agents ex ante. I show that a form of price discrimination can be used when agents are ex ante heterogeneous. Not using such price discrimination is shown to be costly in terms of efficiency in small groups. Small heterogeneous

groups are outperformed by their homogeneous counterparts when price discrimination is not applied. However, this inefficiency in small groups can be eliminated by using price discrimination. The use of price discrimination becomes irrelevant in large groups and heterogeneous groups always outperform their homogeneous counterparts, whether price discrimination is used or not.

Introduction

This thesis is divided into three distinct chapters which analyse the (in)efficiencies of different provision mechanisms and market structures for the provision of different goods. The first two chapters mainly focus on issues of competition, each focusing on a specific market, while the third chapter focuses on the provision of public goods using a voluntary contribution mechanism.

When modelling aggregate allocations that result from the combination of individual decisions of all agents involved, results are driven by many assumptions made within models. Assumptions made on the individual level can change aggregate results greatly. There are two main sets of assumptions that I elaborate on across all three chapters.

The first set comprises of *assumptions about the level heterogeneity of agents*. I allow for high levels of heterogeneity of agents in all three chapters and highlight the effect this has on results. Allowing for a maximum level of heterogeneity appears key to me. It is necessary if one wants to model accurately a world in which agents differ extremely across many dimensions. In the first chapter, I allow for valuations to differ across agents for primary and secondary goods in a model of consumption chains. This is in contrast to most of the literature, which usually only allows for heterogeneity across one dimension. In the second chapter, I compare results about the optimal bias provided by interest and media groups when followers have identical priors to the case in which priors (and in some settings the level of confirmation bias) differ across agents. In the final chapter, I allow for agents to be ex ante heterogeneous as opposed to only ex post heterogeneous and highlight the effect this has on the efficiency level of voluntary contribution mechanisms for public goods.

The second set comprises of *assumptions about the role of information and the level of uncertainty.* In the first chapter, I highlight how equilibrium prices derived under different market structures and their welfare implications depend strongly on the level of uncertainty consumers face about their valuations. In the second chapter, I show how information can develop over time, how private views can be formed and potentially be manipulated, and how all of this depends on competition. In the third chapter, I show how the level of efficiency of voluntary contribution mechanisms for a public good can depend on the level of information individual agents have about other agents.

Asymmetric Complementarity: The Role of Information and Market Structure on Pricing and Welfare

Kimon Doulis
School of Economics
University of Edinburgh

Abstract

Many modern consumption goods are characterised by consumption chains. A primary good is a stand alone product and an essential component. However, consumers can add further components (secondary goods), which are worthless on their own, to the system. The secondary good is, therefore, naturally tied to the primary good, but the reverse does not hold. It seems natural for firms to provide both products, but the question is whether this is desirable, or whether one should be concerned about a single firm having large market power in both markets.

In this paper, I investigate the optimal pricing policies of firms under different market structures, namely when a primary good monopolist is allowed to integrate and have monopolistic power in both markets, when it is faced with competition in the secondary market and when it is prohibited from entering the secondary good market and, hence, has to compete across markets. I show that optimal pricing policies depend strongly on the informational assumptions made on the consumer side of the market, but that a monopolist will never bundle the goods. Overall, an integrated monopolist charges a price in the market for the secondary good that is equal to marginal costs as long as consumers are uncertain about secondary good valuations. An integrated firm, in that scenario, actually appears preferable to trying to introduce competition by prohibiting the firm from integrating. However, no such general result exists when there is no uncertainty about secondary good valuations.

1 Introduction & Literature

A lot of attention has been devoted to the analysis of optimal pricing of multi-product monopolists and how pricing and welfare depend on the level of complementarity or substitutability of the products. Extensive analysis has been provided for the case of complementary products being provided by different firms ever since Cournot (1838). The pricing of multi-product firms can be especially important and have large impacts on overall welfare when the products are strong complements or are part of a consumption chain.

Many modern consumption goods are parts of consumption chains; they are complements that build

on each other. Some goods are base goods and additional goods work as add-ons, upgrades or extensions. Examples for such goods would be operating systems and pieces of software, smartphones and applications, applications and in-app purchases, TV's and DVD players and many others. All these share two properties. Firstly, a good A provides users with utility on its own and is a stand alone product. Secondly, another good B only works in combination with aforementioned good A and may further increase the user's utility. Hence, there are asymmetric complementarities. While product A is improved by the existence of B, good A is completely essential for consumption of good B.

It is the aim of this paper to analyse this setting of asymmetric complementarities. The main focus is going to lie on analysing the optimal pricing strategy of an integrated monopolist providing both products and analyse how optimal prices depend on the informational assumptions made on the consumer side of the market. Furthermore, a second emphasis is going to be put on comparing welfare implications of the integrated firm's optimal pricing policy and comparing welfare results to other market structures as well as the welfare maximising allocation. Finally, I am going to change the assumptions made in order to see how results derived previously rely on them or not. In the beginning of the paper I am only going to consider separate pricing, but I will show that an integrated firm will prefer separate pricing to bundling, i.e. an integrated firm does not make use of tying.[1]

Production and consumption chains

The existing literature has analysed production chains. It has identified a potential double mark-up problem when an intermediate output is produced monopolistically and sold to a monopolistic provider of a final output, who uses the intermediate output as an input. However, such chains are not unique to the production side and it is surprising that there has been more limited focus on consumption chains.

The model of a consumption chain naturally relates to the notion of a potential double mark-up problem in vertically related non-competitive markets. As mentioned, this mark-up problem has been studied on the production side, i.e. in models with intermediate goods, for example by Blair and Kaserman (1978). The literature shows how a monopolistic provider of an intermediate good has a strong incentive to integrate downstream. If firms do not integrate, we are faced with a welfare reducing double mark-up problem. Although the total effect of integration on welfare is

---

[1]The literature on bundling and tying (the difference of which is mainly legal and not a technical one) and the incentives for it is extensive. Although the model presented here seems connected to that literature, mainly because the secondary good is naturally tied to the primary good (but not vice versa), I will show that a monopolist will not make use of bundling, but prefers separate pricing.

not clear-cut, there seems to be more of a case for integration when some goods are intermediate goods. To some extent, I aim to extend this logic to the case of consumption chains.

Tying and bundling

In this paper, two complementary goods are provided either by different firms, by an integrated monopolist or a monopolist in the primary market faced with competition in the secondary market. Especially the latter two suggest some relation to the large literature on tying and bundling. This literature analyses the incentives of multi-product firms with market power to bundle goods and offer them at a single price and to tie goods provided in a competitive market to the good they provide in a monopolistic or oligopolistic market. The early literature on tying and bundling, most prominently Adams and Yellen (1976), highlights the use of bundling as a tool for price discrimination. Bundling can be a profit maximising strategy even in the absence of cost saving. The analysis was later extended by, amongst others, McAfee et al. (1989) who show that the optimality of bundling the goods depends on the correlation of valuations.[2] I will show that bundling is not optimal in the model presented here, because the secondary good is of no value on its own. The nature of the goods considered here is different to the ones considered in the bundling literature. While goods might be complementary in that literature, consumers derive utility from consuming only one of the goods. This is not the case in my model, as the secondary good does not work as a stand alone product. The valuation for the secondary good on its own is zero.

A second strand of the literature on tying, put into a formal model by Whinston (1990), has focused on a second motive, the so called "leverage theory". This literature suggests that a monopolist in market A might find it optimal to tie its good sold monopolistically in market A to a second good provided by it in market B. By doing so the monopolist can leverage its monopoly power into market B, which otherwise would have been more competitive. Once again, I do not formally regard such pricing policies, but I will show that a monopolist in the primary market does not change its pricing policy, whether it is faced with competition in the secondary market or not.

In response to recent anti-trust decisions especially in the software and applications industry (these are typically markets and goods that match our model reasonably well and also potentially provide strong incentives for tying), which tried to reduce the active use of tying, there has been some research, such as Gans (2011), that takes a look into incentives for tying in these markets specifically. In these models, whether tying policies are actually welfare reducing remains largely unclear. Gans (2011) uses a simple model similar to the one used here, in which the primary good is essential, i.e.

---

[2]In general, bundling as a tool of price discrimination seems more appealing when valuations are negatively correlated.

the secondary good is useless on its own. However, he only considers a scenario of a monopolist in the primary market and duopolistic competition in the secondary good market. Furthermore, competition in the secondary market is Hotelling-like, i.e. there exists horizontal product differentiation, which is necessary in his model, because the possibility to bundle would be limited when competition in the secondary good was Bertrand-like. I consider different market structures and focus on how the information of consumers affects pricing in a model without horizontal product differentiation.

In general, the incentives and effects of bundling are reasonably well understood and the literature is extensive. However, in these models both goods typically need to have value on their own, which is not the case in the model presented here. In some sense, the secondary good is already naturally tied to the primary good, but not vice versa, while the question usually is whether a monopolist wants to tie the primary good to a secondary good in the existing literature. By analysing separate pricing, I want to answer the question whether the existence of market power in both markets is necessarily bad in terms of welfare and what our answer to this question depends on. It might be that tying is welfare reducing, but not because it enables the primary good monopolist to leverage its monopoly power into a secondary market, but because of some other intrinsic feature of the tying policy. This would mean we should be concerned about tying, but not about the existence of market power in itself.

Complementary components and systems

There exists a literature analysing pricing and effects of vertical integration when goods are strictly complementary components and together form a system. Typically there exists a variety of all components and a section on the literature focuses on firms' incentives to provide components that are compatible with other components provided by a different firm and analyses potential welfare effects (Matutes and Regibeau, 1988/1992). Farrell and Katz (2000) focus on the incentives of a monopolistic provider of good A to enter a competitive market providing component B and potential effects on innovation in the market for B. Components differ in quality, but consumer preferences are homogeneous. Economides and Salop (1992) extend the idea of Cournot competition to the market for systems and analyse a variety of market structures. However, as with the other papers in this strand of literature, both goods are components but none of them is essential and there is, thus, no notion of a consumption chain. The notion of a consumption chain exists in Ellison (2005), who provides a model of add-on pricing. However, he focuses on a setting in which prices for add-ons are not known by consumers. Firms in his model might want to sell the primary good for low prices and lure consumers into the market with the intention to sell add-ons at high prices.

The consumer might not be aware of this when making the decision to purchase the primary good. In the model presented here, prices are known and binding, but consumers might not be aware of their valuation for the secondary good. The idea that a monopolist providing two complementary goods might find it optimal to sell one at or even below marginal costs is not new (Davis and Murphy, 2000), but I provide a better understanding of this by providing a formal framework and showing under what assumptions this holds when complementarities are asymmetric.

In terms of the setup of the model, the model presented here is related to the literature on patent pools (Lerner and Tirole, 2004; Rey and Tirole, 2013), which is also related to the literature on systems. This literature focuses on the incentives for patent owners to pool patents. Users derive some utility from using/licensing a single patent as well as for a combination of them. The papers show how the degree of complementarity between the patents determines the shape of users' demand functions and, hence, the incentives for pooling. Allowing for the degree of complementarity between the patents and the degree of essentiality of patents to vary, they also show that pools are welfare enhancing if complementarity is large enough. In the model presented here, I focus on a specific setting of asymmetric complementarities. The primary good is fully essential, and I restrict myself to separate pricing, i.e. there is no bundling of goods. I do, however, allow for the valuations of secondary goods as well as primary goods to be heterogeneous and potentially uncertain.

A small literature incorporating the idea of an essential primary good and a complementary second good focusing on separate pricing exists. Economides and Viard (2009) introduce network effects and vertical product differentiation in the secondary good market into this model. However, they do not allow for consumer uncertainty about secondary good valuations and do not provide any general results about pricing or welfare but rather focus on trying to numerically find parameter values which yield identical prices in a pure model with network effects. Cheng and Nahm (2007, 2010) analyse markets in which a secondary good acts as a quality improver to an existing primary good. It matches the model presented here, because the base good is necessary and the good improving the quality is worthless on its own. Cheng and Nahm (2007, 2010) focus on monopolistic competition across markets, i.e. both the primary as well as the secondary goods are provided by independent monopolists. They derive results about equilibrium prices for simultaneous price setting (2007) and sequential price setting (2010). They show that the firms act as if the goods were symmetrically complementary when the secondary good provides only a minor improvement in quality. When it provides a large quality improvement, the goods are priced as if they were independent of each other. However, not only do they focus on a different market structure (I mainly focus on the integrated firm), but they also assume the value of the quality improvement to be the same to all consumers and known.

Overall contributions

This paper combines elements of the different strands of literature. It extends the idea of a double mark-up in production chains to the demand side of the market. In contrast to the bundling literature, the secondary good is already naturally tied to the primary good, meaning that bundling will not be a profit maximising choice for firms. By focusing on a specific case of asymmetric complementarities, namely the primary good being necessary for the secondary good, this paper can, in contrast to existing papers in the literature, analyse the effects of different market structures and different assumptions about the information levels of consumers, while allowing for heterogeneous valuations in both goods. Especially the focus on different assumptions about consumer information is something that has been missing from the current literature on complementary components.

By doing this, the paper helps to provide answers to, amongst others, the following questions:

- How problematic is an integrated firm in a setting of consumption chains?
- Is the extension of market power into secondary goods necessarily an issue to be concerned about?
- How strong is the case for intervention and which of the two markets should the intervention focus on, i.e. is it more crucial to increase competition in the secondary market or in the primary market?
- How do results depend on the information of consumers and is it necessarily always good for consumers to be informed?

Some of the results could be used and applied to antitrust cases such as United States vs. Microsoft Corporation considering illegal tying of Internet Explorer and Windows Media Player or the potential issue of tying Windows and Office. Tying usually is used to extend monopolistic power into a further market. This paper contributes towards answering whether it is the existence of market power in separate markets itself that is problematic. Furthermore, the model provides an alternative explanation as to what determines whether price mark-ups are large for primary or secondary goods.

In the following section, I am going to present the baseline model, in which consumers are not uncertain about their valuations for the secondary good, derive optimal pricing policies, provide a welfare analysis and some comparative statics.[3] In section three, which is the main focus of this paper, I derive results assuming that consumers are uncertain about secondary good valuations.

---

[3]The analysis of this case is going to be less detailed as the literature in this setting is more extensive and I can draw from existing results. In any case, one of the main aims of this paper is to highlight how results depend drastically on the assumptions made about the information of consumers, rather than to necessarily provide quantitative results for each case.

In section four, I will then provide a comparison of the two cases and specifically compare welfare derived under an integrated monopolist in each setting. I will then conclude in the final section.[4]

2 The model without uncertainty

One might think that consumers know all relevant information, including the additional value the secondary good provides to them. This is the assumption made in related literature by Economides and Viard (2009), Davis and Murphy (2000), Gans (2011), Rey and Tirole (2013), and Cheng and Nahm (2007, 2010). Ellison (2005) considers the case in which add-on prices are unobservable at the point consumers decide whether to purchase the base goods, which is arguably similar to uncertainty about valuations. However, rational agents, in his model, in which there exist two goods that differ in quality, can infer the optimal prices. Ellison (2005) shows that the existence of a low quality good enables competing firms to create adverse selection on purpose in order to soften competition. This enables the firms to charge higher add-on prices. The unobservability of prices works in different ways compared to uncertainty in my model. Uncertainty is not used as a tool to create adverse selection and extract consumer surplus in the secondary market, but, as we will see, as a tool to maximise expected consumer surplus and extract it in the primary market.

In this setting, I am going to assume that consumers know all the relevant valuations before making the primary good demand decision. Each agent i knows both realisations a_i and b_i ex ante.[5] I will compare the results of this benchmark case to a case in which secondary good valuations are uncertain at the point consumers decide about primary good demand.

2.1 Consumers, choices and demand

There are two goods: a primary and a secondary good. On top of providing utility from consumption, the primary good, which can be bought at price p_1, is a necessity for secondary good consumption. In other words, the secondary good can only be added to the primary good and cannot be consumed on its own (it provides zero utility on its own). Only consumers who consume the primary good have the option of adding the secondary good at price p_2.

Consumer i's valuation for the primary good is given by a_i, which is the realisation of a random

---

variable $\tilde{a}_i$ distributed according to the density function $f(a_i)$. The density function is identical for all consumers, is strictly positive on the interval $[0, 1]$ and zero otherwise. In order to ensure existence and uniqueness of the equilibria derived in this paper, I will assume that $f(a_i)$ has a decreasing hazard rate, i.e. $\frac{d\frac{f(a_i)}{1-F(a_i)}}{da_i} > 0$.[6] $F(a_i)$ is the corresponding cumulative distribution function. Assume there is a unit mass of consumers, so that the fraction of consumers with some specific valuation a_i is identical to the density $f(a_i)$. The valuation a_i is realised before decisions about primary and secondary good consumptions are made.

The secondary good valuation b_i is distributed independently of the primary good valuation a_i and adds to utility in an additive way. It is the realisation of a random variable $\tilde{b}_i$ which is distributed according to the density function $g(b_i)$. This function is strictly positive on the interval $[0, B]$ and zero otherwise. I will also assume an increasing hazard rate for $g(b_i)$. The parameter B aims to capture how valuable the secondary good is compared to the primary good. I assume $\frac{\partial g(b_i)}{\partial B} < 0 \ \forall \ b_i \leq B$, i.e. the density $g(b_i \leq B)$ decreases when B increases to B', while $g(B' \geq b_i > B)$ increases by definition.

The consumer faces two simple binary decisions of unit demand in both markets. His payoff is simply given by the difference between valuations and prices paid, i.e. it is $a_i - p_1$ if he consumes only the primary good and $a_i + b_i - p_1 - p_2$ if he consumes both goods. Consumers are assumed to be risk neutral.

Starting with the secondary good decision of consumers we can say that a consumer with valuation b_i who has got access to the secondary good, because he purchased the primary good, faces the following simple maximisation problem:

$$\max_{d_2 \in \{0,1\}} (b_i - p_2) * d_2$$

where d_2 is the consumer's discrete choice variable. This leads to $d_2 = 1$ and, therefore, a unit demand if $b_i \geq p_2$.

Aggregating over all consumers and remembering that only primary good demanders have access to the secondary good we get an aggregate demand D_2^c given by:

$$D_2^c(p_1, p_2) = \left(1 - G(p_2)\right) * D_1^c(p_1, p_2)$$

where $D_1(\cdot)$ is primary good demand.

The consumer has to take into consideration that consumption enables him to potentially enjoy an extra benefit of secondary consumption when he makes the primary good consumption decision.

---

[6]This assumption is in line with the literature and corresponds to assumptions about demand elasticities needed to derive a unique profit-maximising monopolistic price. In this paper, the assumption is actually stricter than needed for most results.

Primary good consumption provides the consumer with an option value. If the consumer purchases the primary good, the value of the secondary good to the consumer is given by:

$$V_2(b_i, p_2) = \max_{d_2 \in \{0,1\}} (b_i - p_2) \, d_2.$$

Substituting in the optimal policy rule for d_2, we can say that the secondary good has the following value to consumers:

$$V_2(b_i, p_2) = \begin{cases} b_i - p_2 & \text{for } b_i \geq p_2 \\ 0 & \text{otherwise} \end{cases}.$$

There is no uncertainty about the option value of the primary good, because b_i is known when consumers make their primary good demand decision. All consumers are going to derive a positive individual surplus ex post.[7] The option value is given :

$$V_2(b_i, p_2) = \begin{cases} b_i - p_2 & \text{for } b_i \geq p_2 \\ 0 & \text{otherwise} \end{cases}.$$

Hence, the maximisation problem for the consumer, when faced with prices p_1 and p_2, in the primary stage is given by:

$$\max_{d_1 \in \{0,1\}} \begin{cases} (a_i - p_1 + b_i - p_2) \, d_1 & \text{for } b_i \geq p_2 \\ (a_i - p_1) d_1 & \text{for } b_i < p_2 \end{cases}.$$

where d_1 is the consumer's discrete choice variable.

It follows that $d_1 = 1$ for $a_i \geq p_1$ independent of b_i. I will refer to this group as *direct demanders*. Consumers in this group demand the primary good independent of its additional option value. Direct demand without uncertainty about secondary valuations ($DD_1^c(p_1)$) is independent of b_i and given by:

$$DD_1^c(p_1) = 1 - F(p_1).$$

In addition, $d_1 = 1$ for $p_1 > a_i \geq p_1 + p_2 - b_i$. I will refer to this group as *indirect demanders*. Consumers in this group derive negative benefit from the primary good in isolation. They still demand the primary good, because it provides them with an additional option value. Indirect demand for a consumer depends on the sum of secondary valuation (b_i) and primary valuation (a_i). In order to derive total indirect demand we, therefore, have to integrate over all types. This yields:

$$ID_1^c(p_1, p_2) = \int_{p_2}^{B} \left(\int_{p_1 + p_2 - b_i}^{p_1} f(a_i) da_i \right) g(b_i) db_i$$

$$= (1 - G(p_2)) F(p_1) - \int_{p_2}^{B} F(p_1 + p_2 - b_i) g(b_i) db_i.$$

---

[7]We will see that this is not necessarily the case once we allow for b_i to be uncertain when consumers decide about primary good demand.

Adding the two forms of demands yields a *total primary good demand* of:

$$D_1^c(p_1, p_2) = 1 - F(p_1)G(p_2) - \int_{p_2}^{B} F(p_1 + p_2 - b_i)g(b_i)db_i.$$

Direct demanders with $b_i \geq p_2$ and all indirect demanders are going to demand the secondary good, because $b_i \geq p_2$ is a necessary condition for indirect demand. This means that *total secondary demand* is given by:

$$D_2^c(p_1, p_2) = 1 - G(p_2) - \int_{p_2}^{B} F(p_1 + p_2 - b_i)g(b_i)db_i.$$

Consumer and producer surplus functions, including the fixed point problems defining producer surplus maximising optimal prices (lemmata 1 and 2), for the setting without uncertainty are derived in appendix A.

2.2 Welfare

Proposition 1 (Welfare without uncertainty)**.**

Aggregate welfare is uniquely maximised when prices are equal to marginal costs in each market:

$$p_1^* = c_1$$

$$p_2^* = c_2.$$

The prices p_1^* and p_2^* implement the unique first best allocation.

Proof. Proofs for propositions that do not follow directly from the analysis provided in the main body of the paper can be found in annex B. Supporting lemmata and surplus functions for the analysis without uncertainty about b_i are derived in annex A. □

Some consumers will stop demanding the primary good as soon as $p_1 > c_1$. From a welfare point of view, though, they should demand the good because $a_i > c_1$. For the allocation to be efficient, trade should occur whenever total benefits outweigh total costs and not otherwise. This means that trade in the primary good should occur if $a_i + p_1 \geq p_1 + c_1$ and/or if $a_i + b_i + p_1 + p_2 \geq p_1 + p_2 + c_1 + c_2$. Trade in the secondary market should occur if $b_i + p_2 \geq p_2 + c_2$. These conditions can only be satisfied at the welfare maximising prices identified in the proposition.

2.3 Integrated monopolist

I am assuming constant marginal costs of production of c_1 in the primary market and c_2 in the secondary market. In addition, I assume there to be no fixed costs. A primary good producer's payoff or profit is then given by $(p_1 - c_1) * D_1^c$. The payoff of a secondary good producer is $(p_2 - c_2) * D_2^c$. I provide more detail on the precise nature of the surplus functions in appendix A.

The payoff of the integrated monopolist is simply the sum of individual payoffs. Rewriting our marginal profit conditions, which are derived in appendix A, we can write:

$$\frac{\partial PS^c}{\partial p_2} = \frac{\partial PS^c}{\partial p_1} - (1 - F(p_1))G(p_2) - (p_2 - c_2)(1 - F(p_1))g(p_2) + (p_1 - c_1)f(p_1)G(p_2).$$

This shows that the second first order condition is not a fraction of the first plus an additional negative demand effect (we will see this to be the case in the case with uncertainty about b_i). Price changes always cause negative demand effects in complementary markets. Furthermore, as explained in more detail in annex A1, the distribution of types in the secondary market is skewed towards higher valuations. This makes higher secondary prices more appealing to the monopolist. However, the secondary market is by definition smaller than the primary market. Which effect dominates and in which market the monopolist, therefore, charges a higher mark-up depends on the parameters of the model (on the distributions and especially on the value of the secondary good).

As pointed out by Andriychenko et al (2006) it proves to be very hard to provide general results for prices of separate components charged by a multi-product monopolist when goods are complements. However, by restricting myself to the case in which the primary good is fully essential and the secondary good is not, I am able to provide the following general result.

Proposition 2 (Pricing of the integrated firm in the case without uncertainty).
The integrated firm does not charge a price equal to marginal costs in either market. Instead we can say that $p_2^{PS^c} > c_2$ and that $p_1^{PS^c}$ is larger than a monopolistic price without the existence of a secondary good market, which would be defined by the solution to $p_1 = \frac{1 - F(p_1)}{f(p_1)} + c_1$.

The result is in line with the literature. Economides and Viard (2009) provide numerical results for different parameter settings, but also no further general results. Rey and Tirole (2013) show in a dynamic model of patent pools that coordination, which can be thought of as two firms acting more like an integrated firm, is more likely to occur when complementarity is strong, which is the case in this model. However, they focus on a case in which both goods are equally valuable and essential for each other, while allowing for results to depend on the level of essentiality/complementarity.

15

They provide an extension to this, but focus on the incentives for coordination, rather than on implications on prices. Their paper does, therefore, not provide any further insight into the relative mark-ups on the individual goods.

Intuitively, the distribution of secondary good valuations of consumers left in the market is skewed towards higher valuations, causing the secondary good price to be above marginal costs. The monopolist can also charge a primary good price above a price it would charge if the secondary good did not exist, because most consumers moving onto the secondary market will derive additional benefits. This setting, on its own, provides us with relatively little insight apart from showing that the monopolist derives profits in both markets. Economides and Viard (2009) show, using numerical examples, that the integrated firm will, in general, charge a higher mark-up on secondary goods. As I show, this can be explained by the fact that the consumers who made it in the secondary market are more likely to have high valuations for the secondary good. The firm, therefore, finds it easier to extract surplus in that market. Such pricing strategies can be observed often. Classical examples of this are relatively low mark-ups on operating systems (e.g. Windows) compared to software running on that system (e.g. Office), or the fact that TVs are usually sold at extremely low profit margins while additional components are sold at higher margins. Pricing strategies in other areas, such as cars and add-ons or tablets and applications are more mixed, differ across firms and can actually evolve over time.[8] Finally, I point out that the classical example that is often used for a low base good and very expensive secondary good, the printer and cartridge example, does not match my model. A printer is arguably of no value without a cartridge; the complementarity is not asymmetric.

One might think that the integrated firm would like to bundle the goods, rather than offer separate prices. However, this turns out to not be the case in this model of asymmetric complementarities, because the secondary good is of no benefit on its own.

Proposition 3.

The integrated multi-product monopolist does not make use of bundling. It strictly prefers separate pricing.

---

[8]I will show how the assumption that is placed on the information a consumer has is key to the optimal pricing strategy. Differences across firms, sectors and time in the observed pricing strategy could be explained by different assumption about consumer information, or changes in the level of information consumers have.

2.4 Monopolist and perfect competition

Let us analyse a market structure in which there exists a monopolist in the primary market and in which the secondary good is provided under perfect competition. Introducing a zero profit condition in the secondary market because of perfect competition means that the price in the secondary market will be $p_2^* = c_2$. The secondary price is, therefore, clearly smaller in this market structure.

Faced with perfect competition in the secondary market the primary monopolist charges a price $p_1^{PS_1^c}(p_2^*)$ in the primary market. This price is larger than the price charged by an integrated monopolist in the primary market, because $\frac{\partial p_1}{\partial p_2} < 0$.

The overall effect on welfare is unclear. There exists no general answer to the question which market structure is preferable in terms of welfare. It proves difficult to derive general results, because there is no constant rate of substitution of prices which leaves primary good demand unchanged.[9] The main result here is that these market structures do not result in equivalent allocations.

2.5 Competition across markets

Finally, let us consider a market structure of competition across markets. The goods are provided by separate monopolists in each market. Monopolists in each market set prices simultaneously. We are, thus, looking for the Bertrand-equilibrium of prices.

An equilibrium is a set of prices $p_1^{PS_1^c}$ and $p_2^{PS_2^c}$ such that the reaction functions derived in lemmata 1 and 2 are mutually best responses. This means that equilibrium prices are defined by:

$$p_1^{PS_1^c} = \frac{D_1^c(p_1^{PS_1^c}, p_2^{PS_2^c})}{f(p_1^{PS_1^c})G(p_2^{PS_2^c}) + \int_{p_2^{PS_2^c}}^B f(p_1^{PS_1^c} + p_2^{PS_2^c} - b_i)g(b_i)db_i} + c_1$$

$$p_2^{PS_2^c} = \frac{D_2^c(p_1^{PS_1^c}, p_2^{PS_2^c})}{\int_{p_2^{PS_2^c}}^B f(p_1^{PS_1^c} + p_2^{PS_2^c} - b_i)g(b_i)db_i + (1 - F(p_1^{PS_1^c}))g(p_2^{PS_2^c})} + c_2.$$

This situation has been analysed by Cheng and Nahm (2007) for a model in which the primary good provides a quality of z and the secondary good provides an improvement to the primary good which is identical for all consumers. Rey and Tirole (2013) assume that users differ in terms of the implicit valuation for the overall system. However, there is no difference in the valuation of individual components.[10] Economides and Viard (2009) allow for valuations to be heterogeneous

---

[9]This will be the case for the scenario with uncertainty about secondary good valuations.

[10]They assume that users have to incur a cost that differs across the agents if they decide to adopt the technology, but that adding a second component does not create differential benefits to consumers

for both components, as do I. In Cheng and Nahm (2007), the quality of the primary good upgraded with the secondary good is given by q, with $q > z$. They prove the existence of unique equilibria for some values of z, but show that an equilibrium in pure strategies does not exists for intermediary values of z. Given the complicated nature of the model when q is homogeneous, it seems very unlikely that introducing another dimension of heterogeneity in q, as in the model presented here, would resolve this and provide more clear-cut results.

Although Cheng and Nahm (2007) show that the double mark-up problem is stronger for low z, meaning that welfare under separate pricing is increasing in z, we cannot properly compare these results to the results derived in the case of an integrated monopolist. Once again, as pointed out by Andriychenko et al. (2006), it proves difficult to provide any general results for pricing and welfare in the case without uncertainty, unless one assumes specific parameter values.

2.6 The value of the secondary good

Producer surplus maximising prices under all market structures depend on the parameters of the model. A key parameter of interest is the value of the secondary good. It is, therefore, interesting to see whether any of the allocations implemented by the different market structures get close to an efficient allocation for some values of B.[11]

I am going to restrict my analysis to the case of an integrated firm and a monopolist in the primary market with perfect competition in the secondary market, because there is clearly no tendency for prices to reach marginal costs in competition across markets. In that case, the secondary price is clearly increasing in the value of the secondary good and the primary firms price is also going to increase as the primary firm captures some of the additional value. The more interesting question is whether any of the integrated firm's prices move towards an efficient level.

The unique efficient level of prices $p_i^* = c_i$, which we derived earlier, is independent of the value of secondary good. The integrated monopolist acts as a simple monopolist in the primary market with $B = 0$ and charges prices above marginal costs otherwise as $\frac{\partial p_i^{PS^c}}{\partial B} > 0$. The mark-up above efficient prices is increasing in both markets and the welfare loss is, thus, increasing in the value of the secondary good.

A primary market monopolist faced with perfect competition starts of in the same situation when $B = 0$. There is, by definition, no increase in the secondary good price as B increases, because $p_2^* = c_2$ independent of B. However, $\frac{\partial p_1^{PS_1^c}}{\partial B} > \frac{\partial p_1^{PS^c}}{\partial B} > 0$. The primary good price increases in the

---

[11]We will see that this is the case when b_i is uncertain before primary good consumption.

value of the secondary good more than in the case of an integrated firm.

The allocations of different market structures do not converge to each other with increasing B when b_i is uncertain. The allocations implemented by all market structures considered in this paper do not converge to an efficient allocation.

3 The model with uncertainty

In this scenario, I am going to assume that consumers are uncertain about their secondary good valuation when making their purchasing decision for the primary good. It is only upon consumption of the primary good that b_i is realised. Consumers need to experience the primary good in order to identify how useful a complement the secondary good is. I am not aware that the scenario of uncertainty about secondary good valuations has been considered in models that resemble consumption chains. I have explained that Ellison (2005) allows for unobservable add-on prices. We will see that uncertainty without correlation of valuations will mean that the expected value of the secondary good is identical for all consumers at the point they make their primary good purchasing decisions. One could, therefore, argue that papers like Rey and Tirole (2013) that assume that adding a second component provides the same benefit to consumers are essentially equivalent. However, actual secondary valuations in my model will be heterogeneous when consumers decide whether or not to purchase the secondary good, while they are not in such models. Uncertainty about valuations will have an impact on prices. Having the same expected valuation is not the same as having the same actual valuation.

3.1 Consumers, choices and demand

In this setting, b_i is unknown to consumers at the point of primary good consumption decision making. The option value is, therefore, uncertain and identical for all consumers when they decide whether to buy the primary good or not. Consumers have to factor in the expected value of additional secondary good consumption. Secondary good consumption is purely an option and consumers are going to follow the optimal demand rule derived earlier once b_i is realised. We can, therefore, write the expected option value as:

$$E\left[V_2(p_2)\right] = \int_0^B V_2(b_i, p_2)g(b_i)\, db_i$$
$$= \int_{p_2}^B \left(b_i - p_2\right)g(b_i)\, db_i.$$

A risk neutral consumer with valuation a_i faces the following maximisation problem in the primary good decision stage:

$$\max_{d_1 \in \{0,1\}} \left(a_i + E\left[V_2(p_2)\right] - p_1\right) * d_1$$

which leads to the simple demand rule of $d_1 = 1$ and, therefore, unit demand for $a_i \geq p_1 - E\left[V_2(p_2)\right]$. Once again aggregating over all consumers, we can then derive primary good demand $D_1(\cdot)$ as:

$$D_1(p_1, p_2) = 1 - F\left(p_1 - E\left[V_2(p_2)\right]\right).$$

We can see that primary good demand can again be broken down into two categories (direct and indirect demand), which will enable us to make better welfare comparisons and improve intuition. Firstly, all consumers with $a_i \geq p_1$ demand the primary good directly. Their demand is independent of the existence of the secondary good. Consumers in this group are *direct demanders*. Secondly, consumers with $p_1 > a_i \geq p_1 - E\left[V_2(p_2)\right]$ only demand the primary good because of the additional option value that it provides. These consumers are *indirect demanders*.

We can, thus, write direct demand $\left(DD_1(\cdot)\right)$ and indirect demand $\left(ID_1(\cdot)\right)$ as:

$$DD_1(p_1) = 1 - F(p_1)$$

$$ID_1(p_1, p_2) = F(p_1) - F\left(p_1 - E\left[V_2(p_2)\right]\right).$$

We can see that direct demand under uncertainty is identical to the case without uncertainty $(DD_1(p_1) = DD_1^c(p_1))$. Indirect demand is determined by a single threshold value for a_i above which consumers demand the primary good. In the case without uncertainty, the option value was not a constant. Two agents with identical a_i could make different demand decisions; the sum of valuations determined demand. This is not the case when b_i is uncertain, because the option value of the primary good is identical for all consumers.

In general, I will refer to the cut-off primary good valuation, above which consumers demand the primary good, as a^{min}, i.e. $a^{min}(p_1, p_2) = p_1 - E\left[V_2(p_2)\right]$. Furthermore, I will refer to $E\left[V_2(p_2)\right]$ as EV_2 at times.

The relations and sizes of the different forms of primary demand are plotted in their simplest form in figure 1, in which primary good valuations are plotted against willingness to pay. The blue line depicts demand triggered purely by the primary good, while the red line shows demand taking the option value into account. The shape of the blue line is determined by the distribution of primary good valuations. In figure 1 it is a straight line, because a_i is assumed to be uniformly distributed. With b_i being uncertain, the option value is independent of the type a_i and, thus, just a parallel upward shift of the blue line. The magnitude of the shift depends on B, p_2 and the distributional form $G(\cdot)$.
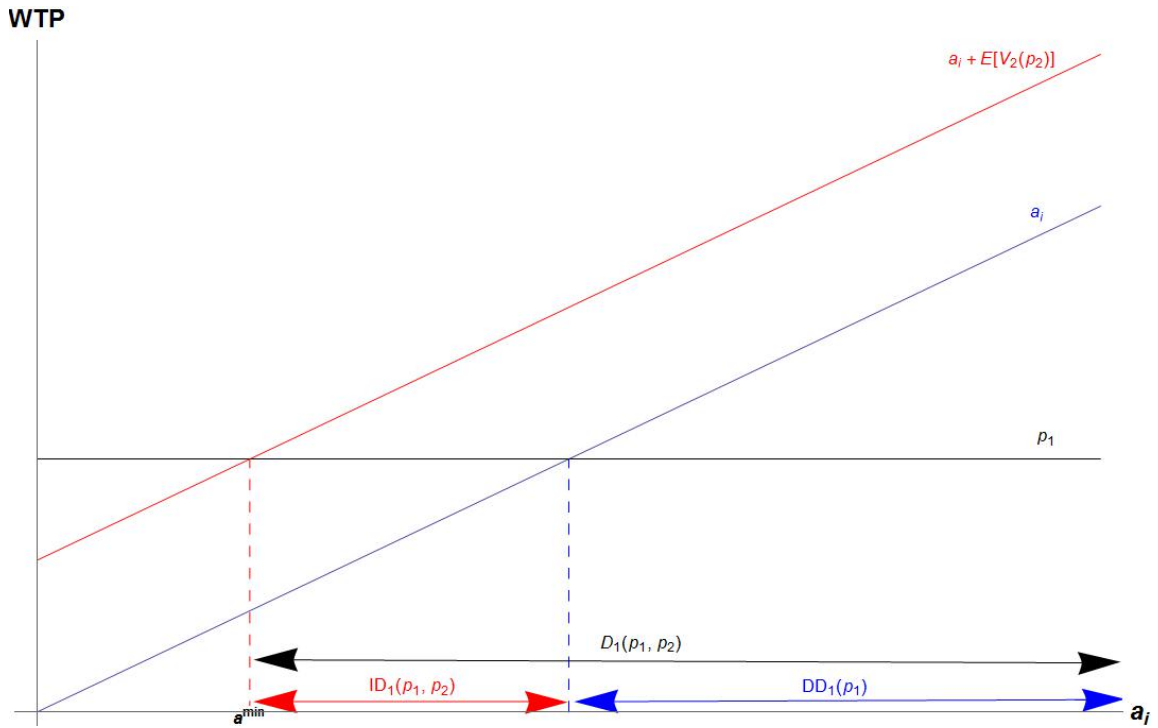
Figure 1: Demand in the primary market

Figure 2 illustrates the effect of an increase of the primary price from p_1 to p_1^N on the types of demand in the primary market. The price increase causes a clear reduction in direct demand ($\frac{\partial DD_1}{\partial p_1} = -f(p_1)$). It causes flow from direct into indirect demand and out of indirect demand (into non-demand). The total effect on indirect demand is given by: $\frac{\partial ID_1}{\partial p_1} = f(p_1) - f(a^{min}(p_1, p_2))$. For the uniform distribution depicted, the two flows will cancel out exactly and the size of indirect demand, therefore, remains unchanged. The change in a^{min} is, thus, identical in size to the change in p_1 (for the uniform distribution depicted here). An exception would be when $a^{min} < 0$. In that case there is full demand and, therefore, only a flow into indirect demand.
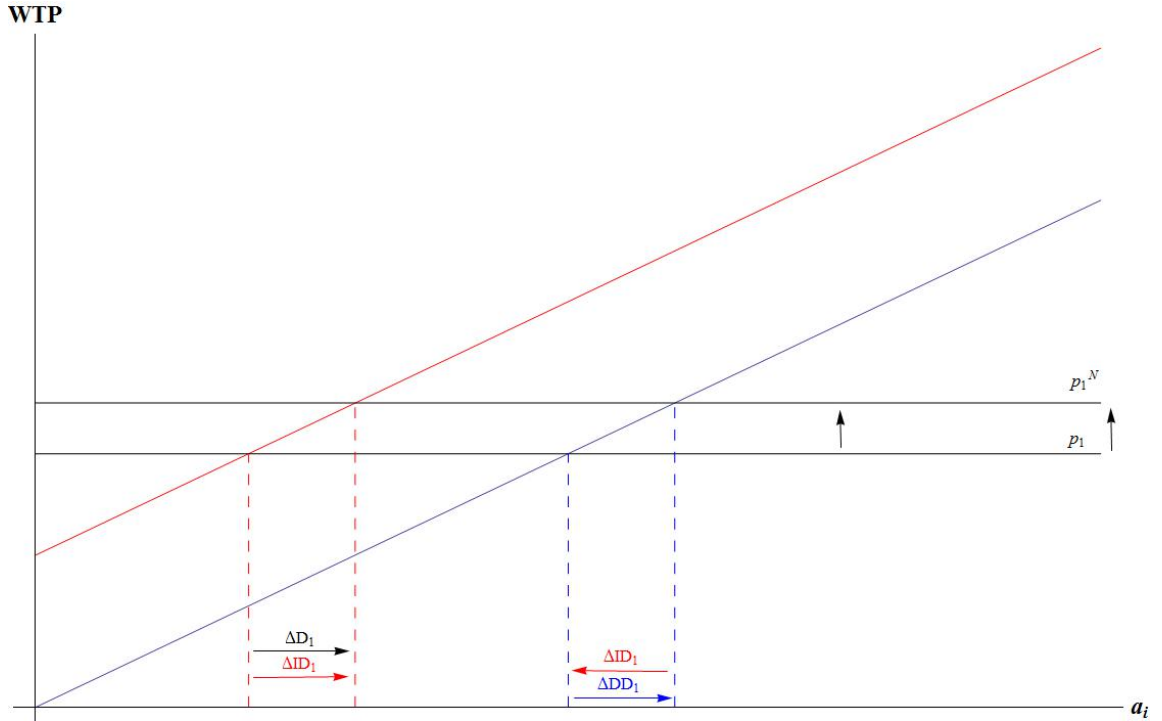
21

WTP

Figure 2: Effect of a change in p_1 on primary good demand

The change in p_1 affects the secondary market by reducing the mass of consumers that moves into it. This decreases secondary demand $\left(\frac{\partial D_2}{\partial p_1} = -f(a^{min})(1 - G(p_2))\right)$.

Figures 3 and 4 illustrate the effects of an increase of the secondary price from p_2 to p_2^N on the types of demand. Such a price change leaves direct demand unchanged $\left(\frac{\partial DD_1}{\partial p_2} = 0\right)$ and only causes a flow out of indirect demand (into non-demand) by decreasing the option value $\left(\frac{\partial ID_1}{\partial p_2} = -f(a^{min})(-\frac{\partial EV_2}{\partial p_2})\right)$. Since $\frac{\partial a^{min}}{\partial p_2} = -\frac{\partial E[V_2(p_2)]}{\partial p_2} = 1 - G(p_2, B)$, changes in secondary price cause a change in a^{min} smaller than the price change itself. However, changes in p_2 cause two negative demand effects in the secondary market. As was the case for primary price changes, secondary good price changes cause a decrease in the mass that moves into the secondary market. Secondary good price changes cause a second effect by changing the fraction of primary good demanders that demand the secondary good (shown by the blue arrow).
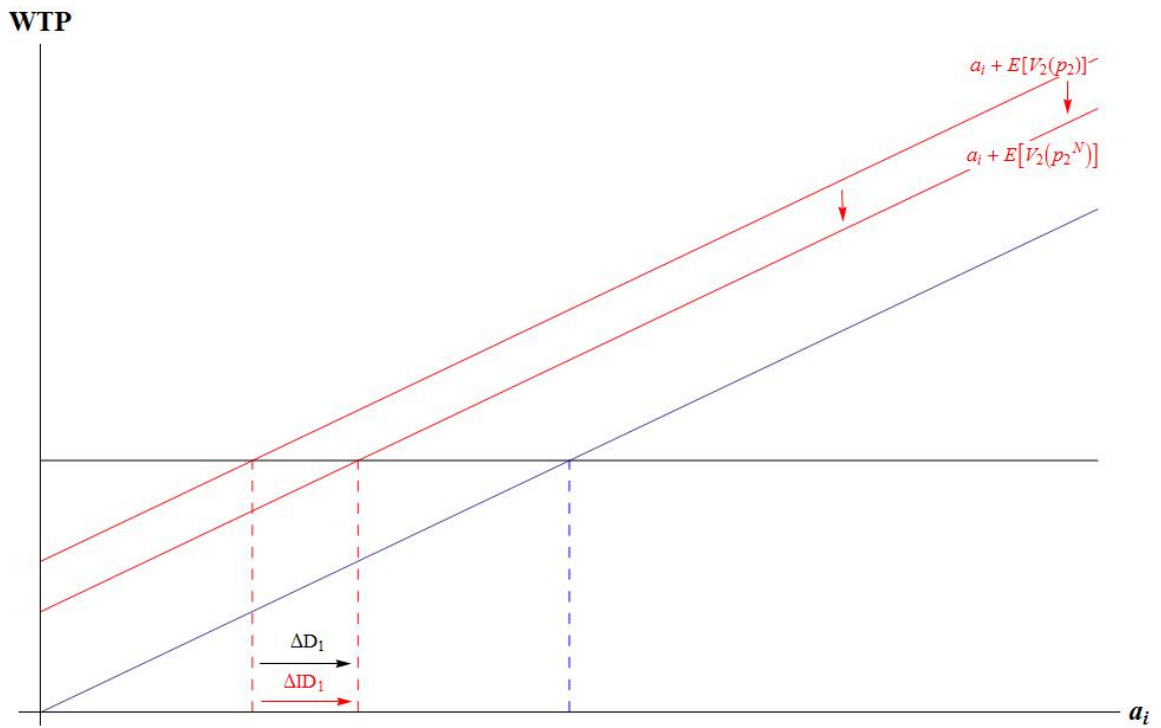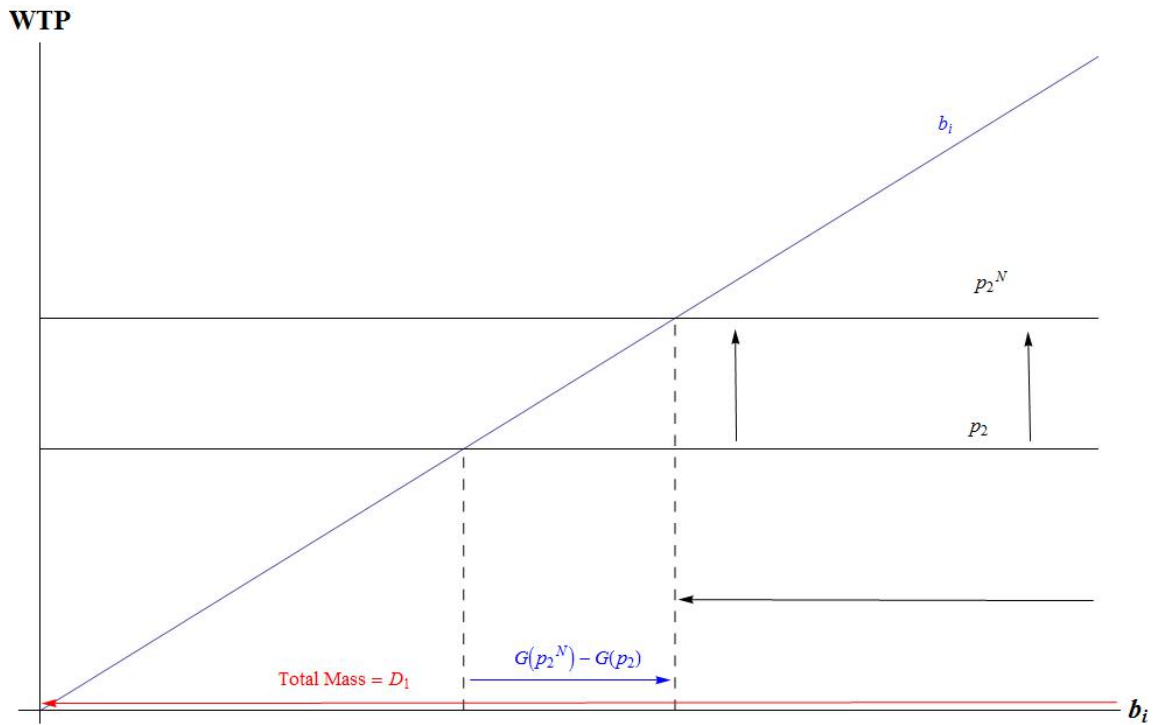
Figure 3: Effect of a change in p_2 on primary good demand

Figure 4: Effect of a change in p_2 on secondary good demand

3.2 Consumer surplus

Using the demand rules derived and aggregating over all consumers we derive total ex post consumer surplus derived from primary good consumption as:

$$CS_1(p_1, p_2) = \int_{a^{min}}^{1} (a_i - p_1)f(a_i)\, da_i.$$

This can be broken down into primary good consumer surplus of direct demanders and indirect demanders:

$$CS_{1,DD}(p_1) = \int_{p_1}^{1} (a_i - p_1)f(a_i)\, da_i$$

$$CS_{1,ID}(p_1, p_2) = \int_{a^{min}}^{p_1} (a_i - p_1)f(a_i)\, da_i$$

where the latter is negative, because indirect demanders only demand the primary good, because the option value offsets the negative benefit of the primary good. It can already be seen that this will cause some consumers to regret their decision ex post. There will be *'regretted demand'* ex post for indirect demanders whose realisation b_i turns out to be small.

All primary good demanders move on to the secondary good market. Of these, each consumer demanding the secondary good derives a surplus of $b_i - p_2$ from it. Aggregating individual surpluses we get our measure for consumer surplus derived from secondary good consumption:

$$CS_2(p_1, p_2) = D_1(p_1, p_2) * \left[\int_{p_2}^{B} (b_i - p_2)g(b_i, B)\, db_i \right]$$
$$= \left(1 - F(a^{min})\right) * E\left[V_2(p_2)\right]$$

where the average benefit derived from secondary good consumption ex post is just given by the ex ante valuation $E\left[V_2(p_2)\right]$. This is unsurprising, because risk neutrality and the assumption of a unit mass of consumers ensure that ex post shares of consumers with specific valuations are identical to the ex ante probabilities.

Finally, total consumer surplus, $CS(\cdot)$, is simply the sum of the three separate measures. That is:

$$CS(p_1, p_2) = CS_{1,DD}(p_1) + CS_{1,ID}(p_1, p_2) + CS_2(p_1, p_2).$$

One of the main aims of this paper is to analyse and break down welfare effects of specific price structures in consumption chain markets. In order to do this, we have to understand how individual consumer surpluses, as well as overall consumer surplus, react to price changes.

After basic simplification, the effect of primary good price changes is given by:

$$\frac{\partial CS}{\partial p_1} = F(a^{min}) - 1.$$

The negative effect of primary price changes on aggregate consumer surplus is decreasing in a^{min} and, hence, increasing in aggregate demand, because it is a simple direct effect, i.e. each consumers derives a smaller surplus when p_1 increases. Apart from this direct effect, there are two additional demand effects. Some consumers switch from being direct demanders to indirect demanders. This, however, leaves overall surplus unchanged. Furthermore, a mass of $f(a^{min})$ stops demanding the primary good at all. Since these are marginal consumers for which the loss of primary good consumption is exactly outweighed by the expected benefit of the secondary good, the increase in consumers surplus by indirect demanders in the primary market, $f(a^{min})E\left[V_2(p_2)\right]$, is exactly cancelled by the decrease in secondary good consumer surplus $,-f(a^{min})E\left[V_2(p_2)\right]$, leaving overall consumer surplus unchanged.

Deriving the effect of secondary good price changes on consumer surplus yields:
$$\frac{\partial CS}{\partial p_2} = (F(a^{min}) - 1)(1 - G(p_2, B)).$$
The secondary good price causes a smaller change in consumer surplus overall, because it only works indirectly via a decrease in the option value. It also causes additional demand effects, but, as in the case of changes in the primary good price, this does not affect overall consumer surplus, because positive and negative effects cancel out.

In summary, primary good price increases cause a large fall in primary good demand caused by a fall in direct demand, while secondary good price increases cause a smaller fall in primary good demand due to a change in indirect demand. A marginal change in the secondary good price causes a smaller change in consumer surplus because it affects a smaller market.

Furthermore, increases in p_2 decrease the fraction of consumers who regret their purchasing decision ex post, because increases in p_2 decrease the primary good's option value. There is less 'regretted demand' from the consumers' perspective. Increases in p_1 leave the size of indirect demand and, therefore, the fraction of ex post 'regretted demand' fairly unchanged by creating a flow into indirect demand from direct demand and out of indirect demand into non-demand, where the overall effect depends on the distributional forms. Primary price increases only create a flow into indirect demand only when in full demand $(a^{min} \leq 0)$.

Since the consumer surplus function is continuous and decreasing in both prices, it is obvious that:
$$\arg\max_{p_1, p_2} CS(p_1, p_2) = (0, 0),$$
while, as shown, the decrease is stronger in an increase in the primary good price. The analysis is summarised in proposition 4.

Proposition 4 (Consumer surplus)**.**

- *There will be 'regretted' demand ex post from the consumers' perspective. Higher secondary good prices eliminate large 'regretted' demand ex post, while it is largely unaffected by primary good prices.*

- *Increases in the secondary good price cause consumer surplus to decrease by less than increases in the primary good price.*

The idea of 'regretted demand' does not exist in the current literature on consumption chains as valuations are not assumed to be uncertain. One should not that this result is not unique to consumption chains. Clearly, decisions of some agents are regretted ex post whenever uncertainty is added to any kind of model.

Figure 5 shows the consumer surplus functions for three numerical examples. Primary and secondary valuations are uniformly distributed in all examples. The functions differ in the value of the secondary good. The first graph depicts a case of $B = \frac{1}{2}$, the second $B = 1$ and the third $B = 2$. The red area highlights the area of full demand, in which $a^{min} \leq 0$. The graphs aim to convey the idea that secondary good price increases are less damaging for consumer surplus.
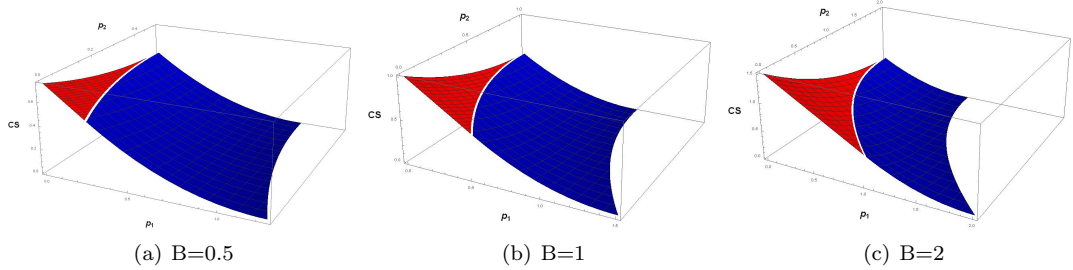

(a) B=0.5 (b) B=1 (c) B=2

Figure 5: Consumer surplus functions

3.3 Producer surplus

The separate producer surplus functions are given by:

$$PS_1(p_1, p_2) = (p_1 - c_1)D_1(p_1, p_2)$$

$$PS_2(p_1, p_2) = (p_2 - c_2)D_2(p_1, p_2).$$

Total producer surplus is just their sum:

$$PS(p_1, p_2) = PS_1(p_1, p_2) + PS_2(p_1, p_2).^{12}$$

26

I am interested in how producer surpluses/profits are affected by different price structures and how they react to price changes. I need to identify producer surplus maximising prices to establish and compare the allocations implemented by the different market structures I consider. Firstly, have a look at the effect of primary good price changes on the producer surplus functions.

Lemma 1.

- *The unique primary good price that maximises primary good producer surplus must be in the following range:*

$$p_1^{PS_1} \in \left[E\left[V_2(B, p_2)\right], 1 + E\left[V_2(B, p_2)\right]\right].$$

  It is the solution to the following fixed point problem:
$$p_1^{PS_1} = \frac{1 - F(a^{min}(p_1^{PS_1}, p_2, B))}{f(a^{min}(p_1^{PS_1}, p_2, B))} + c_1.^{13}$$

- *All primary good prices in the following set maximise secondary producer surplus:*

$$p_1^{PS_2} \in [0, E\left[V_2(B, p_2)\right]).$$

- *The unique primary good price that maximises overall producer surplus must lie in the following interval:*

$$p_1^{PS} \in \left[E\left[V_2(B, p_2)\right], 1 + E\left[V_2(B, p_2)\right]\right)$$

  It is defined by:
$$p_1^{PS} = \frac{1 - F(a^{min}(p_1^{PS}, p_2, B))}{f(a^{min}(p_1^{PS}, p_2, B))} - (p_2 - c_2)(1 - G(p_2, B)) + c_1.$$

Proof of lemma 1. The marginal profit functions are given by:
$$\frac{\partial PS_1}{\partial p_1} = \left(1 - F(a^{min})\right) - (p_1 - c_1)f(a^{min})$$
$$\frac{\partial PS_2}{\partial p_1} = (p_2 - c_2)f(a^{min})\left(G(p_2, B) - 1\right)$$
$$\frac{\partial PS}{\partial p_1} = \left(1 - F(a^{min})\right) - (p_1 - c_1)f(a^{min}) + (p_2 - c_2)f(a^{min})\left(G(p_2, B) - 1\right).$$

---

[12]Producer surplus functions have a kink at $p_1 = EV_2(\cdot)$, since $f(0) > 0$ and, therefore, $F(a^{min}(\cdot))$ takes a discrete jump up at $a^{min}(\cdot) = 0$. Although non-smooth, right and left limits are identical and the producer surplus functions are, thus, continuous in prices.

[13]The inverse hazard function (the fraction) is nothing else in this setting but the reciprocal of the own price elasticity of primary good demand, ϵ_{1,p_1}. The above fixed point problems could be written as a standard monopolistic pricing rule:
$$p_1^{PS_1} = \frac{1}{\epsilon_{1,p_1}(a^{min}(p_1^{PS_1}, p_2))} + c_1.$$
As usual, the more elastic demand is, the smaller will be the price mark-up, while the own price elasticity of primary good demand depends on both prices. The same holds for all other producer surplus maximising prices. Only the elasticities differ and prices maximising overall producer surplus take into account cross-price elasticities.

A change in p_1 causes a positive price mark-up effect on PS_1, while causing a negative demand effect represented by the second part of the marginal profit function. When in full demand, $p_1 \leq E\left[V_2(B, p_2)\right]$, a price change causes no negative demand effect, because $a^{min} < 0$ and $f(a^{min}) = 0 \forall a^{min} < 0$. The first marginal profits function is, therefore, strictly positive in this case. Price increases up to $p_1 = E\left[V_2(B, p_2)\right]$ are, therefore, strictly increasing primary good producer surplus. The optimal price itself is determined by the solution to $\frac{\partial PS_1}{\partial p_1} = 0$. We know this solution will exist and be unique (for non-trivial levels for c_1) as long as our decreasing hazard rate assumption holds. The first bullet point follows directly from this.

The effect of a change in p_1 on secondary good producer surplus is straightforward. The negative demand effect still exists as long as we are not in full demand, while there is no positive price mark-up effect. The marginal profit function, given our assumptions, is zero as long as $a^{min} < 0$ and negative otherwise. Hence, changes in p_1 have no effect on PS_2 for $p_1 < E\left[V_2(B, p_2)\right]$, and a strictly negative effect otherwise. The second bullet point follows directly from this.

Since the overall marginal profit function is strictly positive for full demand, $\frac{\partial PS}{\partial p_1}\Big|_{0 \leq p_1 < EV_2} = 1$, we know that it cannot be that $a^{min} < 0$. The price itself is once again determined by $\frac{\partial PS}{\partial p_1} = 0$, which leads to the fixed point problem presented in the proposition.

Unique solutions to all the fixed-point problems defining optimal prices in this and the following lemma exist, because surplus functions are non-smooth but continuous; marginal surpluses are either strictly positive at the lower bound of the domain and negative at the upper bound, or negative in the entire domain. The marginal profit function $\frac{\partial PS_1}{\partial p_1}$ has a jump discontinuity (jumps down) at $p_1 = EV_2$, but is weakly positive at zero prices and constant up to $p_1 = EV_2$. It is then decreasing, assuming a increasing hazard rate function, and strictly negative at the upper bound of the support. Hence, the optimum is given at the jump point if $\frac{\partial PS_1}{\partial p_1}$ jumps down to a value below zero at $p_1 = EV_2$. If the function jumps down to a value greater than zero, the first order condition must, based on the intermediate value theorem, be satisfied somewhere in $[EV_2, 1 + EV_2]$. Assuming that the hazard rate function is increasing, the second marginal profit function $\left(\frac{\partial PS_2}{\partial p_2}\right)$ has the same properties with regards to the jump point and is also strictly decreasing to the right of the jump point and negative at the upper bound of the support $p_2 = B$. The only difference is that it is also strictly decreasing for all prices below full demand. Hence, based on the intermediate value theorem, we know that the first order condition must be satisfied somewhere in $[0, B]$. Since overall producer surplus is just the sum of individual surpluses, the same logic applies for the existence of a unique solution to the fixed-point problem defining p_1^{PS}. $\qquad \square$

In comparison to $p_1^{PS_1}$, which is determined solely by the own price elasticity, p_1^{PS} has to also account for a change in profits derived in the secondary market. This means that:

$$p_1^{PS} < p_1^{PS_1} \text{ for } p_2 > c_2$$

$$p_1^{PS} > p_1^{PS_1} \text{ for } p_2 < c_2$$

$$p_1^{PS} = p_1^{PS_1} \text{ for } p_2 = c_2.$$

Now have a look at the effects of a secondary good price changes on producer surplus functions:

Lemma 2.

- *The unique primary producer surplus maximising secondary good lies in:*

$$p_2^{PS_1} \in [0, Max[0, p_2^{PS_1} \text{ s.t. } p_1 - EV_2(p_2^{PS_1}) \leq 0]].$$

- *The secondary producer maximising secondary good price lies in:*

$$p_2^{PS_2} \in [0, B].$$

It is defined by:

$$p_2^{PS_2} = \frac{(1 - F(a^{min}))(1 - G(p_2^{PS_2}, B))}{(1 - F(a^{min}))g(p_2^{PS_2}, B) + f(a^{min})(1 - G(p_2^{PS_2}, B))^2} + c_2.$$

- *The total producer surplus maximising secondary good price lies in:*

$$p_2^{PS} \in [0, B]$$

It is defined by:

$$p_2^{PS} = \frac{(1 - F(a^{min}))(1 - G(p_2^{PS}, B)) - (p_1 - c_1)f(a^{min})(1 - G(p_2^{PS}))}{(1 - F(a^{min}))g(p_2^{PS}) + f(a^{min})(1 - G(p_2^{PS}))^2} + c_2.$$

Proof of lemma 2. Marginal profit functions are given by:

$$\frac{\partial PS_1}{\partial p_2} = (p_1 - c_1)f(a^{min})(G(p_2, B) - 1)$$

$$\frac{\partial PS_2}{\partial p_2} = (1 - F(a^{min}))(1 - G(p_2, B))$$

$$- (p_2 - c_2)\left((1 - F(a^{min}))g(p_2, B) + f(a^{min})(1 - G(p_2, B))^2\right)$$

$$\frac{\partial PS}{\partial p_2} = \frac{\partial PS_1}{\partial p_2} + \frac{\partial PS_2}{\partial p_2}$$

A change in the secondary price has a clear weakly negative demand effect on the primary producer surplus, just as a change in the primary price had on the secondary producer surplus. The effect is non-existent in the case of full demand, $p_1 < E\left[V_2(B, p_2)\right]$, and is otherwise strictly negative. The marginal profit function of the primary good producer is, therefore, zero for $a^{min} < 0$ and negative

29

otherwise. The first bullet of the proposition follows directly from that. Note that the negative demand effect of the secondary price is smaller in size than the effect of a change in p_1 on PS_2, because it only works via decreasing the option value.

Changes in p_2 on PS_2 have, as can be seen from the marginal profit function, three effects. Two of the effects are intuitively similar, but quantitatively different, to the effects of changes in p_1 on PS_1. The first effect is positive and caused by an increase in the profit margin. It is shown by the first term in the marginal profit function. Secondly, there is a negative demand effect. An increase in p_2 decreases the option value. This means that less consumers move into the secondary good market as shown by the third term. Finally, changes in p_2 have a second negative demand effect that did not exist when analysing the effect of primary good price changes on PS_1. A smaller fraction out of all consumers that move into the secondary market demands the secondary good when p_2 is increased, while this fraction is independent of p_1. Hence, a change in p_2 still effects PS_2 through a negative demand effect in full demand, which was not the case for a change of p_1 on PS_1 in full demand. We can, therefore, not provide a reduced interval for $p_2^{PS_2}$. As shown in the second bullet, it is the solution to $\frac{\partial PS_2}{\partial p_2} = 0$ and can take all non-trivial values.

The same logic explained in the previous paragraph applies to the proof of the third bullet of the lemma. $\qquad\square$

As with p_1^{PS}, p_2^{PS} in comparison with $p_2^{PS_2}$ has to take into consideration the additional negative effect it causes in the primary market. This means:

$$p_2^{PS} < p_2^{PS_2} \text{ for } p_1 > c_1$$

$$p_2^{PS} > p_2^{PS_2} \text{ for } p_1 < c_1$$

$$p_2^{PS} = p_2^{PS_2} \text{ for } p_1 = c_1.$$

Figure 6 depicts primary, secondary and overall producer surplus functions, using the same numeric examples as before ($a_i \sim U[0,1]$ in all cases and $b_i \sim U[0,0.5]$, $b_i \sim U[0,1]$ and $b_i \sim U[0,2]$ respectively) for zero marginal costs.
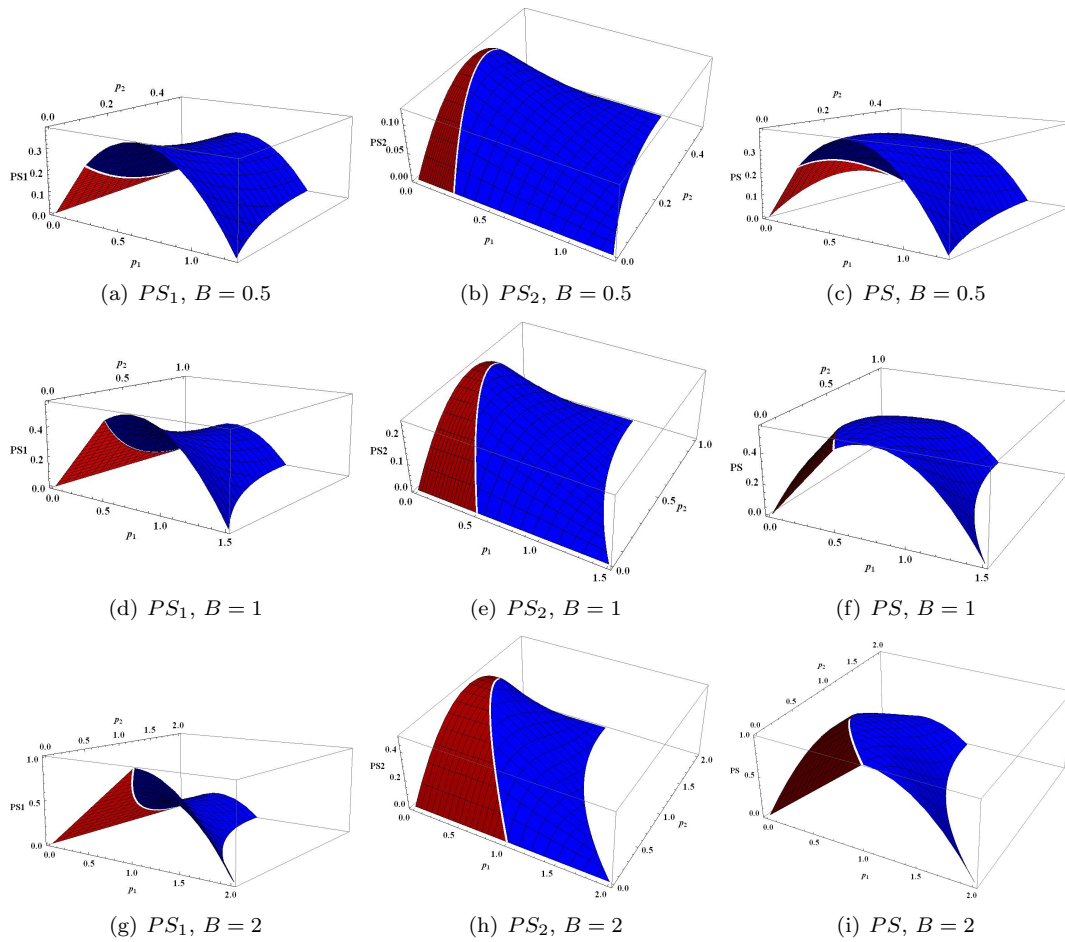
(a) PS_1, $B = 0.5$ (b) PS_2, $B = 0.5$ (c) PS, $B = 0.5$

(d) PS_1, $B = 1$ (e) PS_2, $B = 1$ (f) PS, $B = 1$

(g) PS_1, $B = 2$ (h) PS_2, $B = 2$ (i) PS, $B = 2$

Figure 6: Producer surplus functions

3.4 Welfare

Our basic measure for aggregate welfare is, once again, the sum of consumer and producer surpluses:

$$W(p_1, p_2) = CS(p_1, p_2) + PS(p_1, p_2).$$

We can write the effect of primary good price changes on aggregate welfare as:

$$\frac{\partial W}{\partial p_1} = -(p_1 - c_1)f(a^{min}) - (p_2 - c_2)f(a^{min})(1 - G(p_2)).$$

The direct effect of a higher or lower price cancels out, because the price paid is a pure transfer from consumers to producers. We are left with the simplified effect shown above, which is purely about changes in demand triggered by the price change. A primary price increase has no negative effect on welfare as long as $p_1 < EV_2$ or, in other words, $a^{min} < 0$. It only has a pure redistributive

effect from consumers to producers.

After some basic rearrangements, we can write the effect of changes in secondary good prices on aggregate welfare as:

$$\frac{\partial W}{\partial p_2} = \frac{\partial W}{\partial p_1}(1 - G(p_2)) - (p_2 - c_2)(1 - F(a^{min}))g(p_2).$$

We can say that welfare is decreasing in the secondary good price for $a^{min} < 0$. This was not the case for the primary good price. Which price has a stronger negative effect an aggregate welfare when there is no full demand depends on whether the second demand effect in the case of a change in p_2 outweighs the smaller first demand effect or not. For example, if $g(p_2)$ was large, a change in p_2 would cause a large demand effect in the secondary market and, hence, a larger change on aggregate welfare. The same logic would hold for high levels of p_2, because the additional negative demand effect would be very strong.

In general, high levels of p_2 decrease the fraction of consumers with ex post 'regretted demand'. As shown earlier, the consumer surplus is decreasing in the secondary good price. 'Regretted demand', though, might still be good in terms of welfare. Demand that leads to negative consumer surplus is not necessarily bad in terms of welfare.

Proposition 5 (Welfare maximising prices with uncertain secondary good valuations).
When welfare is measured as the sum of consumer and producer surpluses:[14]

- *the welfare maximising secondary good price is uniquely determined by $p_2^* = c_2$; and*
- *the welfare maximising primary good price p_1^* is not unique for $c_1 \leq EV_2(c_2)$. In this case, there exists a set of prices, $p_1^* \in [0, EV_2(c_2)]$, that ensures the welfare optimum.*

This set of prices exists, because there is full demand up to some primary good price level. Primary good price changes do not cause a demand effect in that case; in other words, demand is fully inelastic. Only demand effects actually cause inefficiencies compared to first best, with direct price effects being purely redistributive. While welfare is equal for all prices contained in $p_1^*(c_1, c_2)$, prices at the lower end of the set cause a high consumer surplus and low producer surplus and prices at the upper end cause the opposite. The results in the proposition are closely related to the existence of 'regretted demand'. No such results have been derived specifically in the literature dealing with consumption chains. The result does, though, not depend crucially on consumption chains. A set of efficient prices could exist in any model with uncertainty when welfare is measured in a way done in this model.

---

[14]In general, one could consider $W = \alpha CS + (1 - \alpha)PS$. The results here hold for $\alpha = \frac{1}{2}$. Although this might seem like a very specific parameter and a policy maker might conceivably choose $\alpha > \frac{1}{2}$, I would argue that equal weights are standard in a general equilibrium model.

All that is relevant in terms of efficiency is that a transaction occurs when total benefits of both market sides (here $a_i + p_1$) outweigh total costs (here $p_1 + c_1$). Welfare maximisation is ensured and first best implemented as long as all consumers with $a_i \geq c_1$ purchase the primary and all consumers with $b_i \geq c_2$ purchase the secondary good. When the primary good has an additional option value that is uncertain, a price above marginal costs does not change this efficiency result as long as it is not too high. Demand that is ex post regretted from the consumers' perspective is not necessarily regretted in terms of efficiency. Prices at the upper end of the set of efficient prices cause a larger fraction of indirect demand and, hence, more ex post regretted demand from the consumer perspective. This causes a big redistribution from consumers to producers. This implies that high prices, for example, for tablets are not necessarily a welfare concern when consumers are uncertain about the additional value that applications and further add-ons can provide.

The proposition, thus, suggests that the existence of supernormal profits is not necessarily a sign of an inefficient market allocation in markets with consumption chains. While it is true that such prices decrease consumer surplus, they have no effect on welfare as long as they do not cause demand effects. In terms of policy, this suggests that any decision about whether prices in a market are excessive should not necessarily be based on a comparison to marginal costs, but on a analysis of the behaviour of demand. The economic case for taking action is limited if there are no demand effects, even if prices are high. This is unless one was to apply a higher weight to consumer surplus than to producer surplus. While this might sound like a result that could be used for defending large redistribution from consumers to producers in such markets, it also implies something very different. It shows that firms can derive supernormal profits when necessary, for example in markets with high costs for R & D, without necessarily harming welfare.

Figure 7 plots aggregate welfare functions for the same three cases as used so far, namely zero marginal costs and uniform distributions with $B = 0.5$, $B = 1$ and $B = 2$. Welfare is constant in the primary good price under full demand, and there exists a set of efficient primary good prices.
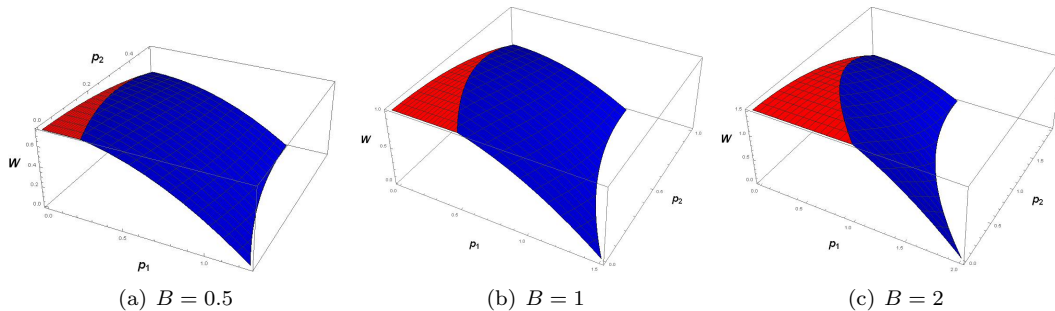

(a) $B = 0.5$ (b) $B = 1$ (c) $B = 2$

Figure 7: Welfare function

3.5 Integrated monopolist

I will now analyse how different market structures compare to the first best benchmark. Firstly, I will consider an integrated monopolist that provides both goods without competitors in either market. Its objective function is given by $PS(p_1, p_2)$. The prices solving this simple maximisation problem, p_1^{PS} and p_2^{PS}, are the solutions to the basic first order conditions and are defined in the lemmata 1 and 2, which also include proof of uniqueness and existence of these prices.[15]

Proposition 6 (Integrated monopolist with uncertain secondary good valuations)**.**

- *The integrated monopolist will charge $p_2^{PS} = c_2$. When consumers are uncertain about their secondary good valuation and the primary good acts as an entry barrier.*

- *The primary good is sold at a mark-up above marginal costs. This mark-up is decreasing in primary good marginal costs, increasing in the option value and decreasing in secondary good marginal cost.*

Intuitively, the option value is identical for all consumers and the firm can extract the ex ante option value by charging it on top of the primary good price.

Unsurprisingly, the optimal primary good price is increasing in own marginal costs. The optimal primary good price is decreasing in secondary marginal costs and increasing in the value of the secondary good, because the option value depends negatively on secondary marginal costs and positively on the value of the secondary good.

The idea that an integrated firm providing complementary goods might charge a price equal to marginal costs for one good is not new. Davis and Murphy (2000) explain that the monopolist might even charge prices below marginal costs for one good in a basic model with linear demand and two complementary goods.[16] Davis and Murphy (2000) use a model without uncertainty; their result is driven by the relative value (and production costs) of the goods. My result is independent of these parameters and is purely driven by uncertainty. One can observe very different pricing strategies in the software, application and video game sectors. While some firms charge low initial prices to lure consumers into consumption and then charge high prices for add-ons to the basic game, software or app (which is opposite to this result), the reverse can be observed as well, i.e. a high price for the basic software, app or game, while updates and add-ons can be downloaded

---

[15]In essence, assuming an increasing hazard function ensured existence and uniqueness even though profit functions are not smooth.

[16]I extend the model in appendix C and show that charging a price below marginal costs in the secondary market is optimal when valuations are negatively correlated.

free of charge. One can also observe a change in pricing strategies in some sectors, which could be explained by a change in consumer information over time. For example, modern TVs are nowadays sold at production costs. Using the optimal pricing of the integrated firm derived in this paper, one could argue that consumers have become much more familiar with additional components that add value to the base good. The uncertainty about the valuation for secondary goods that complement the base good has decreased over time. In response, the integrated firm would increase the mark-up on secondary goods and reduce the mark-up on the primary good.

A technical proof for the proposition is provided in appendix B. I am going to provide a more intuitive proof in the main body of the paper to provide further intuition for this result about the producer surplus maximising combination of prices. Assume that an integrated monopolist currently charges prices p_2' and p_1'. Consider a simple deviation in which the monopolist changes the secondary good price p_2' by an amount ϵ while adjusting the primary good price p_1' such that the primary good demand remains unchanged.

We can treat

$$\frac{\frac{\partial D_1(\cdot)}{\partial p_2}}{\frac{\partial D_1(\cdot)}{\partial p_1}} = G(p_2) - 1$$

as the rate of substitution between prices, holding primary good demand constant. It is worth mentioning that this rate is identical to $\frac{\partial EV_2}{\partial p_2}$. An increase in the secondary good price leads to a decrease in the option value. The primary good price has to be decreased by exactly the change in the option value in order to hold demand constant. In other words, when p_2 changes by ϵ, p_1 can change by $(G(p_2) - 1)\epsilon$ without effecting primary good demand.

How will such a change of prices affect producer surplus?

$$\Delta PS = \Delta p_1 D_1 + (p_2' + \epsilon - c_2)\Delta D_2 + \epsilon D_2$$

Substituting in the ratio of prices that holds $D_1(\cdot)$ constant and the change of secondary good demand, we get:

$$\Delta PS = (G(p_2) - 1)\epsilon D_1 - (p_2' + \epsilon - c_2)g(p_2)\epsilon D_1 + (1 - G(p_2))\epsilon D_1.$$

This leaves us with the following condition for a beneficial (profit-increasing) deviation:

$$\Delta PS \geq 0 \quad \text{for} \quad \begin{cases} -(p_2' + \epsilon - c_2)g(p_2) \geq 0 & \text{if } \epsilon > 0 \\ -(p_2' + \epsilon - c_2)g(p_2) \leq 0 & \text{if } \epsilon < 0 \end{cases}.$$

Increases in the secondary good price are producer surplus increasing as long as $p_2' < c_2$. Decreases in p_2 are producer surplus increasing as long as $p_2' > c_2$. This means that there exist no beneficial deviation only if $p_2' = c_2$.

As long as prices are above marginal costs in the secondary market, an integrated monopolist can reduce the secondary good price while increasing the primary good price in a way that leaves demand in the primary market unchanged. The revenue effects in both markets cancel out, because consumers are risk neutral and the option value is identical to all consumers. However, the integrated monopolist sells the secondary good profitably to a few more customers causing profits to increase.

One might think that the integrated monopolist might prefer bundling rather than charging individual pricing. This, as was the case in our benchmark case without uncertainty about b_i, proves not to be the case.

Proposition 7 (No bundling with uncertain secondary good valuations)**.**
The integrated monopolist never prefers to offer the primary and secondary good as a bundle at a single bundle price. Separate pricing is (weakly) preferable.

The secondary good is intrinsically useless on its own, which is what differentiates this model from the standard bundling literature. If the good has zero value on its own, then why should bundling be preferable?

3.6 Monopolist and perfect competition

Next, I will consider the case where a monopolist provides the primary good, while the secondary good is supplied in a market with perfect competition. With no fixed costs in the secondary market and a zero profit condition, we can say that the price in the secondary market is going to equal marginal costs. This means that the simple monopolist's first order condition in the primary market is identical to the integrated monopolist's one. The only difference being that $p_2 = c_2$ is once an optimal choice and once taken as given (the simple monopolist would obviously, by lemma 2, prefer $p_2 = 0$). We can sum this up in the following proposition.

Proposition 8 (Primary monopolist and competition in the secondary market)**.**

$$p_1^{PS_1} = p_1^{PS} \quad for \quad p_2 = c_2.$$

Prices in both markets are identical, whether the firm is an integrated monopolist in both markets or solely a monopolist in the primary market faced with perfect competition in the secondary market. Both market structures result in the same allocation.

This might not be very surprising, but in general one could expect things to go either way. On the one hand, the integrated firm has market power in both markets and one would, therefore, expect higher aggregate prices causing lower overall welfare. On the other hand, the integrated firm takes into consideration negative externalities caused by high prices in market i on demand in market j. This gives the integrated monopolist an incentive to decrease overall prices. This is beneficial to welfare as high primary good prices inhibit consumers from further consumption in the secondary market. As it turns out, neither of the effects dominates and so both market structures lead to exactly the same allocation.

The proposition suggests that increasing competition in the secondary market has no effects on prices and surpluses as long as the primary market is monopolistic. The allocation is completely independent of the market structure in the secondary good market. This holds as long as consumers are uncertain about secondary good valuations.

3.7   Competition across markets

Finally, we consider Bertrand competition across markets, but not within markets.[17] Primary and secondary goods are provided monopolistically, but they are provided by different firms. I am going to focus on simultaneous price setting rather than sequential price setting and am, therefore, looking for the Bertrand rather than the Stackelberg equilibrium.

By maximising individual producer surplus functions with respect to own prices, we can derive the optimal reaction function of a monopolist in market i, $p_i^{PS_i}(p_j)$, and the equilibrium induced by monopolists in each market and competition across markets. The reaction function of the primary producer $p_1^{PS_1}(p_2)$ (as derived in lemma 2) is defined by $\frac{\partial PS_1}{\partial p_1} = 0$, or:

$$p_1^{PS_1} = \frac{1 - F(a^{min}(p_1^{PS_1}, p_2))}{f(a^{min}(p_1^{PS_1}, p_2))} + c_1.$$

We cannot explicitly solve for the reaction function without specifying the distributions. We can, nevertheless, show how the optimal price $p_1^{PS_1}$ reacts to changes in the secondary price, i.e. derive its slope:

$$\frac{\partial p_1^{PS_1}}{\partial p_2} = \left(\frac{f(\cdot)^2}{2f(\cdot)^2 + (1 - F(\cdot))f'(\cdot)} - 1 \right) (1 - G(p_2))$$

This means that the reaction function is downward sloping with $0 \geq \frac{\partial p_1^{PS_1}}{\partial p_2} \geq -(1 - G(p_2))$. This is unsurprising, because the goods are complements.

---

[17]This is not classical Bertrand competition within a market. I would still refer to it as at least Bertrand-like, because firms compete in prices.

The reaction function of the secondary producer $p_2^{PS_2}(p_1)$ (derived in lemma 1) is defined by $\frac{\partial PS_2}{\partial p_2} = 0$, or:

$$p_2^{PS_2} = \frac{(1 - F(a^{min}(p_1, p_2^{PS_2}, B)))(1 - G(p_2^{PS_2}))}{(1 - F(a^{min}(p_1, p_2^{PS_2}, B)))g(p_2^{PS_2}) + f(a^{min}(p_1, p_2^{PS_2}, B))(1 - G(p_2^{PS_2}))^2} + c_2.$$

Once again, we cannot explicitly solve for the reaction function without specifying distribution functions. It is, however, easy to verify that its slope is weakly negative with $0 \geq \frac{\partial p_2^{PS_2}}{\partial p_1} \geq -1$.

Proposition 9 (Prices and welfare under competition across markets)**.**

Introducing competition across markets:

- *increases secondary good prices and decreases primary good prices ($p_1^{PS_1} < p_1^{PS}$ and $p_2^{PS_2} > p_2^{PS}$); and*

- *decreases aggregate consumer surplus and producer surplus and, hence, also welfare.*

The result that the monopolistic setting can improve consumer surplus and overall welfare in a setting of consumption chains is explained by Andriychenko at al. (2006). Andriychenko, though, considers the scenario in which the monopolist charges one price for the overall system (which is not its optimal choice in the model here). He shows that the implemented allocation dominates the allocation implemented by separate pricing of competing firms in terms of consumer surplus and welfare. This general result is also derived by Davis and Murphy (2000). Finally, Rey and Tirole (2013) explain that patent pools can increase consumer surplus and welfare when patents are strong complements. I have established that these results hold especially when valuations for secondary goods are uncertain.

Figure 8 show optimal prices of an integrated firm, welfare maximising prices and, finally, prices under Bertrand competition depending on marginal costs. Primary prices are plotted in blue and secondary prices are depicted in red. The market structure of a primary monopolist and competition in the secondary market is left out, because it, as shown, is equivalent to the integrated monopolist. All results are derived with both valuations being standard uniformly distributed.

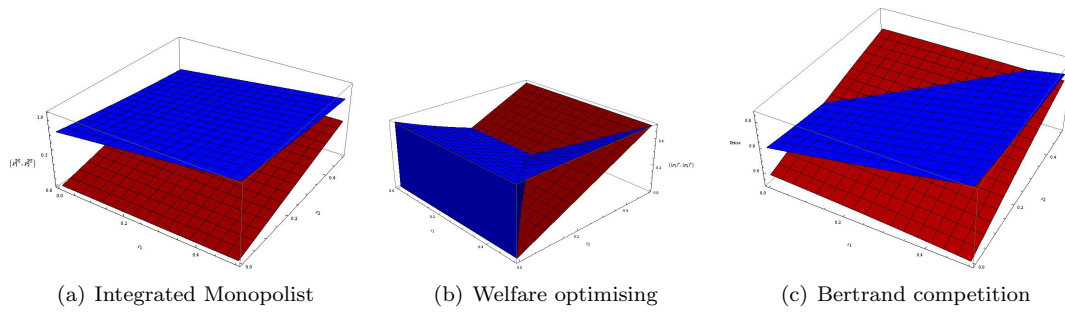(a) Integrated Monopolist (b) Welfare optimising (c) Bertrand competition

Figure 8: Equilibrium and Welfare optimising prices

Figure 9 shows equilibrium prices under the different market structures and first best efficient prices. The red plane depicts prices set by the integrated firm, the blue plane depicts prices derived under Bertrand competition and the green plane depicts the set of efficient prices). The red plane is missing in the graph depicting secondary prices, because the integrated firm's secondary price is identical to the welfare maximising, first best efficient price. It is, therefore, represented by the green plane as well. All results are derived with both valuations being standard uniformly distributed.
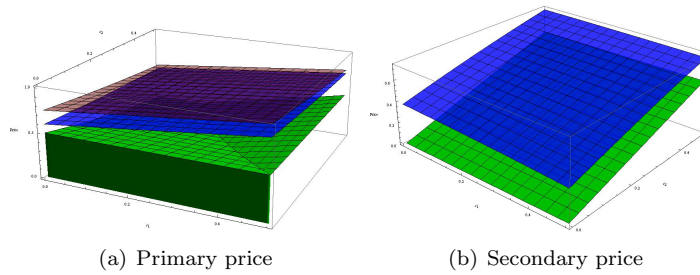

(a) Primary price (b) Secondary price

Figure 9: Price comparisons

3.8   Value of the secondary good

We have shown that the welfare maximising set of prices is not necessarily unique and that the integrated firm provides an allocation that is preferred by consumers and firms to the case of independent monopolists in each market.[18] We have also seen that increasing competition in the secondary market did not create a preferable allocation.[19] I am now going to analyse how the value

---

[18]This result is very much in line with the more famous example of production chains. In production chains, integration is usually preferred by all participants of the market, because it avoids the double mark-up problem.

[19]It seems obvious that competition in the primary market would be beneficial, because this is where the integrated firm charges above marginal costs. However, perfect competition in both markets is always welfare maximising. What if we are currently constrained to a scenario in which some firms have considerable market power or need to derive

of the secondary good affects these results. This analysis will provide an answer to the question of how the results depend on whether the secondary good is just a small add-on or a substantial extra component.

Proposition 10 (Value of the secondary good)**.**

- *The optimal primary good price charged by an integrated monopolist, p_1^{PS}, converges to EV_2 as B increases and remains there for further increases of B.*

- *Welfare derived under an integrated monopolist, taking into account its optimal pricing policy, converges to the first best welfare level.*

- *No such convergence in prices or welfare levels exists for competition across markets.*

The integrated firm always charges a first best price in the secondary market $(p_2^{PS} = p_2^* = c_2)$. It extracts all its surplus in the primary market. The optimal primary good price it charges increases with the value of the secondary good, but the increase in the price is smaller than the increase in the option value of the primary good itself. This means that, for a sufficiently large B, there will be full demand and the allocation implemented under the integrated firm is welfare maximising.

Figure 10 shows how the primary price charged by an integrated firm converges to the upper bound of efficient primary good prices derived earlier for the case of standard uniformly distributed valuations. The graph depicts both the set of efficient prices (remember that welfare is constant for all prices in the set, while consumer surplus is high at the lower end and low at the upper bound and vice versa for producer surplus) and p_1^{PS} as a function of the value of the secondary good. Since $p_2^{PS} = p_2^* = c_2 \ \forall \ B$ the allocation is efficient for all $B \geq 2$ in this example. In general, it gets closer to the efficient allocation as the secondary good becomes more valuable.[20]

---

supernormal profits, because of high costs for R&D etc.? How strong is the case to increase competition in this scenario and in which of the two markets is competition more important?

[20]This can actually be seen on the earlier graph showing the producer surplus function for $B = 2$ (figure 6(i)) since the maximum occurs exactly on the boundary between the red and blue area, i.e. the firm charges a price that induces full demand.
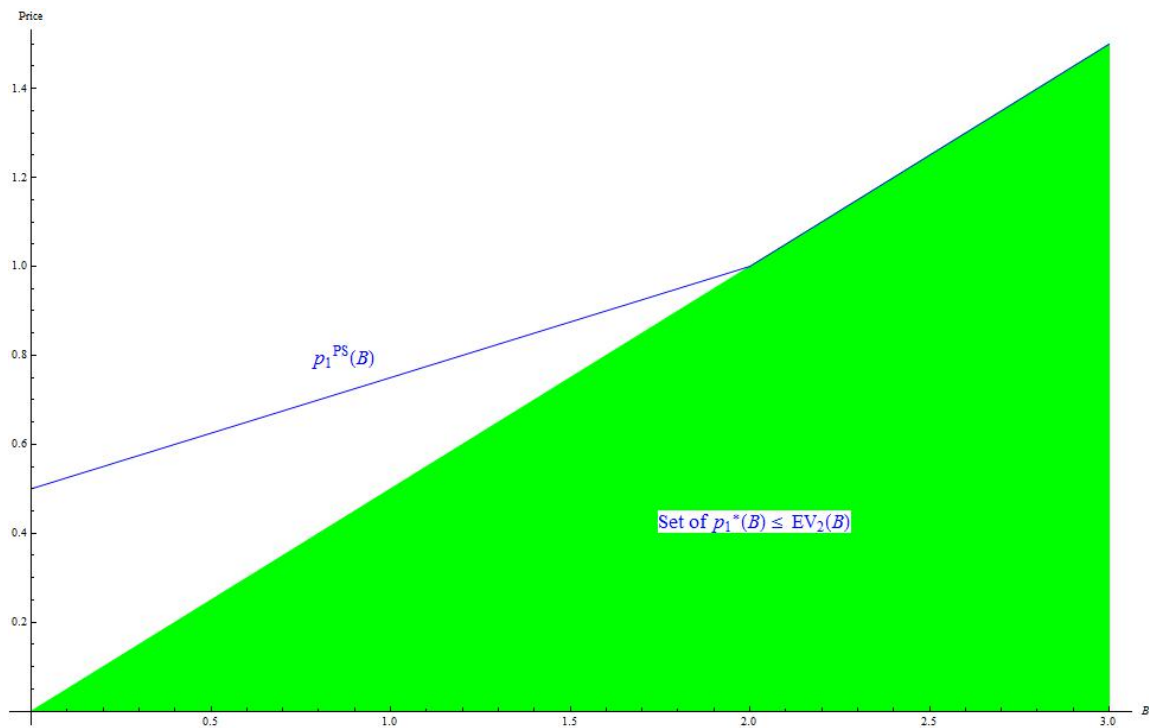
Figure 10: Efficient p_1 and optimal p_1 charged by an integrated firm

4 Comparing welfare with and without uncertainty

We have derived three main results for the case of uncertainty about valuations of goods further down the consumption chain:

- There exists a set of efficient primary good prices.

- The allocation under the integrated firm is identical to a simple monopolist in the primary market with perfect competition in the secondary market.

- The integrated firm charges welfare maximising prices for large B.

Contrasting these results to the case without uncertainty, we can say that all of them depend on the uncertainty of secondary valuations.

An additional question to ask is then which scenario is actually preferable. Can uncertainty actually be good? It seems natural that consumers would prefer a world without uncertainty, because there is no 'regretted demand' ex post. From the consumers' perspective only agents who should demand the good do so. However, whether the case without uncertainty is preferable for consumers is not

41

that clear, because one has to take into account that the equilibrium prices depend on whether we are in a world with or without uncertainty.

On the other side, producers find it much easier to extract consumer surplus under uncertainty. As seen, they extract the surplus in the secondary market by maximising the option value (which is identical for all consumers) and charging it as a mark-up on the primary good price.

Given that a firm has market power in the primary market in the model presented and, thus, that at least some prices are going to be above marginal costs, uncertainty can be a good thing in terms of our simple welfare measure. Usually, the problem in terms of efficiency is that trade does not occur, although there are gains of trade, when prices are above marginal costs. This is the case when consumers do know their secondary valuations. With uncertainty, it is true that some consumers will ex post regret their demand, but in terms of efficiency the trade is efficient even ex post.

This leads to the result that uncertainty about valuations might be a good thing when we are constrained to a world of firms with market power and, therefore, prices above marginal costs for primary goods. This is the case especially when marginal costs are low and option values high, because the option value increases demand; may it be demand that is regretted by consumers ex post.

These results are depicted below (in figures 11-13) for the case of zero marginal costs and standard uniformly distributed primary and secondary good valuations. Surpluses are evaluated at total producer surplus maximising prices, i.e. I am comparing welfare derived under an integrated firm. Surpluses for each type, a_i, are distributed uniformly between the two red lines, which depict the bounds of the surplus. The blue line is the average surplus. Hence, the area under the blue line depicts aggregate consumer surplus. We can also see that some consumers will derive a negative surplus under uncertainty.
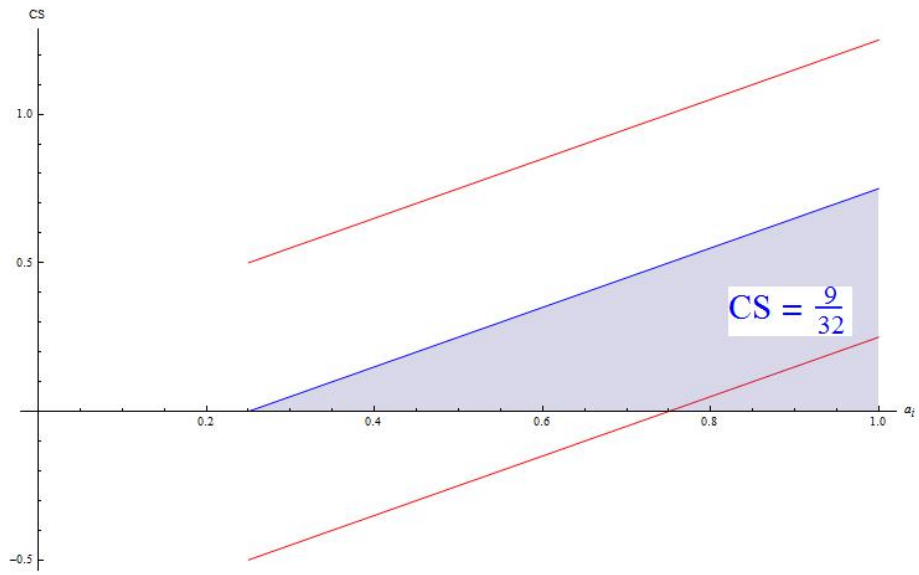
Figure 11: Consumer surplus under uncertainty
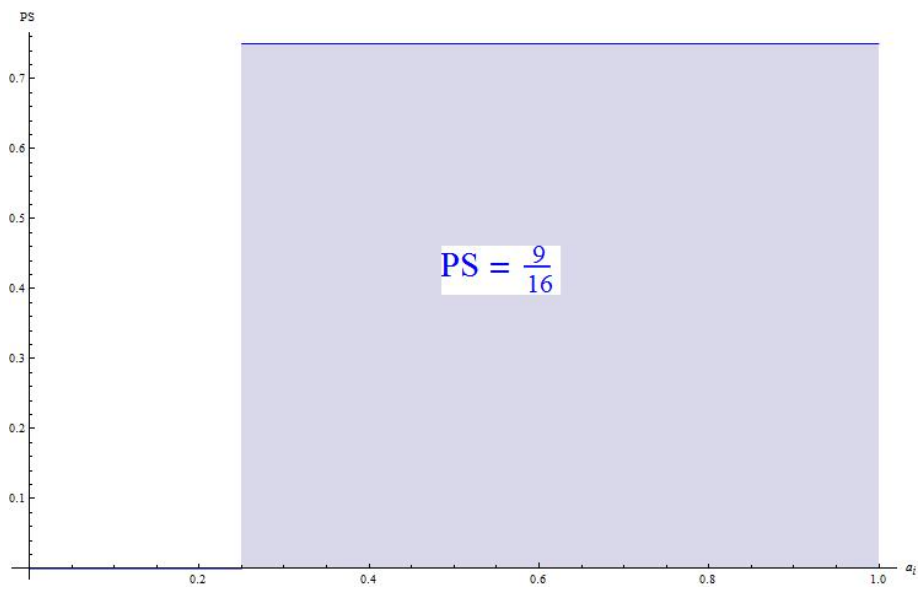
Figure 12: Producer surplus under uncertainty
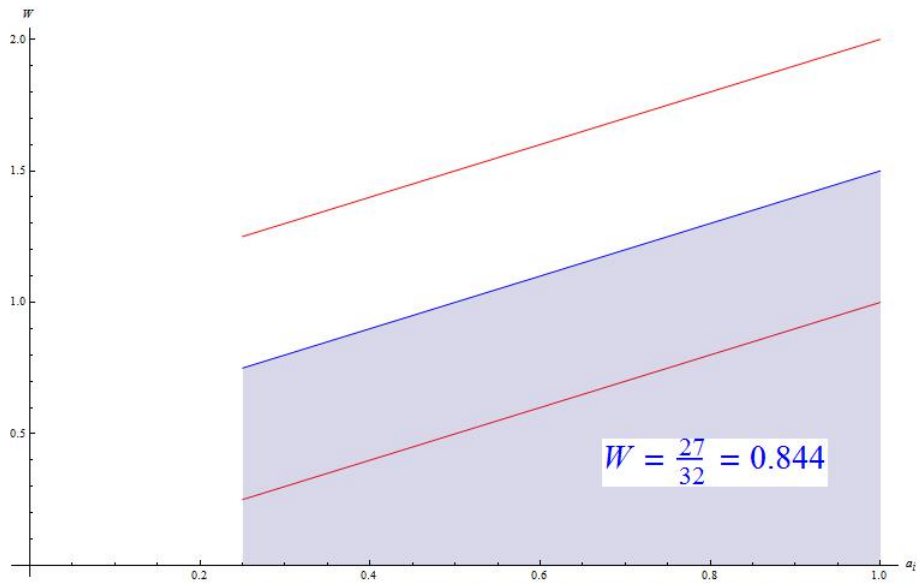
Figure 13: Welfare under uncertainty

Figures 14-16 depict the case of certainty, which is again evaluated at producer surplus maximising prices. This time no consumer derives a negative benefit, while the producer surplus for all direct demanders has an element of uncertainty, because the firm charges a price above marginal costs in the secondary market. As usual the blue line depicts averages, while red lines are upper and lower bounds of the distribution.
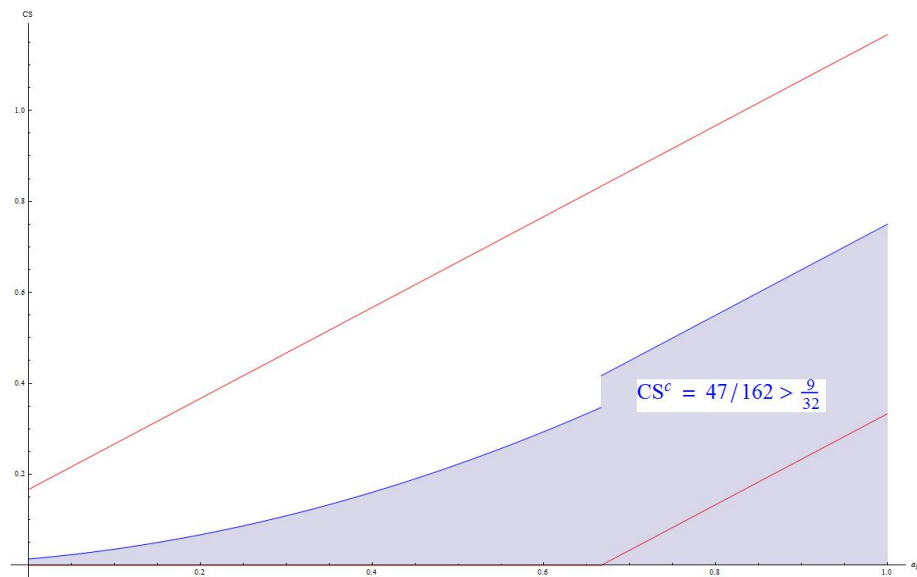
Figure 14: Consumer surplus under certainty
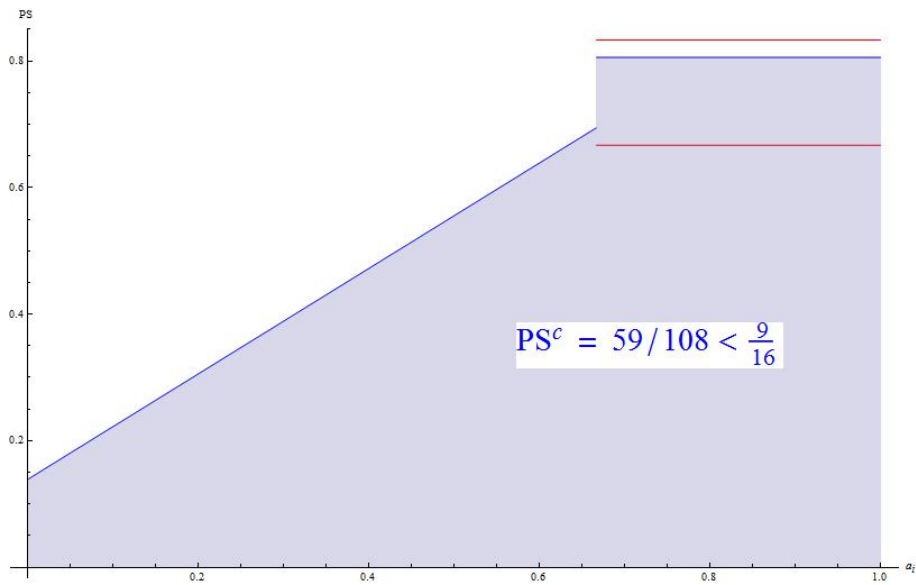
Figure 15: Producer surplus under certainty

$$\mathrm{PS}^c = 59/108 < \frac{9}{16}$$
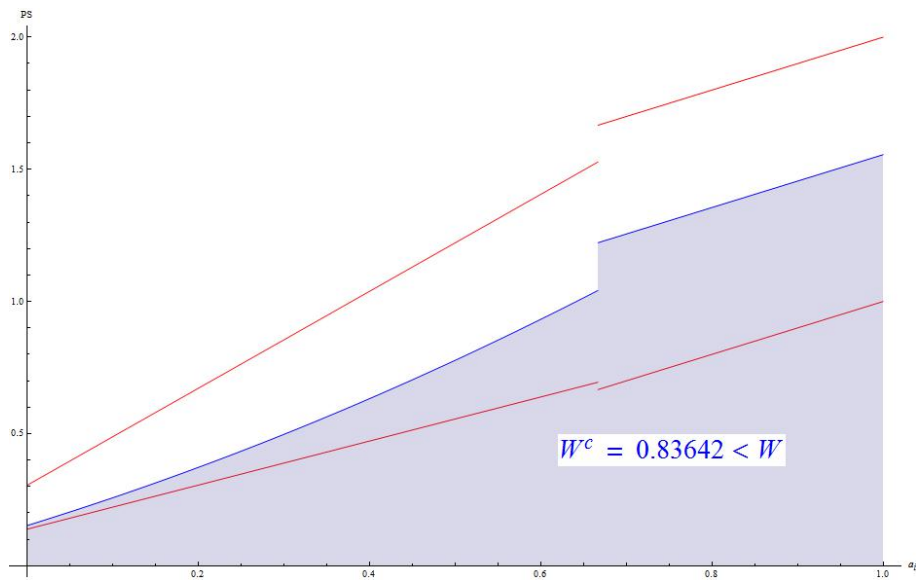
Figure 16: Welfare under certainty

$$W^c = 0.83642 < W$$

It can be seen that consumers, in this example, overall prefer the world of certainty, while the integrated firm clearly prefers uncertainty, because uncertainty enables the firm to extract a large part of the consumer surplus. Since uncertainty triggers a larger demand and full demand is efficient in this numerical example, welfare derived under uncertainty is actually larger.

This shows that uncertainty about valuations for goods further down the consumption chain can be a good thing in terms of aggregate welfare. I have, however, also pointed out the potential distributionary concerns of the allocation under uncertainty. It also remains questionable whether one would want to assign equal weights to consumer and producer surplus in a social welfare function. Note that this result is not specific to the case of consumption chains, but potentially holds if markets are above marginal costs in any market.

5 Conclusion

When one good is essential for the consumption of further goods asymmetric complementarities arise. One can model markets like this as consumption chains. We have seen that optimal prices of a multi-product monopolist (an integrated firm providing the primary and the secondary good) as well as prices of monopolists competing across markets depend strongly on whether consumers are informed or face uncertainty about their secondary good valuations. The same can be said about welfare measures for different market structures.

An integrated firm derives all profits in the primary market when primary and secondary good valuations are uncorrelated and consumers are uncertain about secondary valuations. The firm charges marginal costs in the secondary market. The prices and allocation implemented by the integrated firm is equivalent to a simple monopolist in the primary market faced with perfect competition in the secondary market. Furthermore, for a basic welfare measure, there exists an entire set of welfare-maximising primary good prices. Welfare is constant within this set, while consumer surplus is decreasing in the price and producer surplus increasing. We have also seen that the integrated firm's optimal primary good price gets closer to the upper end of this set as the secondary good becomes valuable. For some value of B, i.e. if the secondary good is very valuable, the integrated firm charges a welfare maximising price (the highest price within the efficient set). There is no such convergence towards efficiency for a market structure with competition across, but not within, markets. Finally, welfare under the integrated firm is strictly larger than under monopolistic competition across markets and the gap between the two is increasing in the value of the secondary good. This seems to suggest that, when constrained to a scenario where a firm has large market power in the primary good market, it might be counterproductive to try to increase competition by not allowing it to integrate. Introducing competition in the secondary market has no effect and prohibiting the monopolist from entering the secondary good market decreases welfare and consumer surplus.

However, all these results depend strongly on uninformed consumers. Without uncertainty about secondary valuations, the integrated firm faces a skewed distribution of secondary good valuations of primary good demanders. Charging secondary good prices equal to marginal costs is not the integrated firm's optimal price policy in the secondary market anymore. Hence, there is no equivalence to a simple monopolist and perfect competition in the secondary market. The integrated firm instead derives profits in both markets and there is no convergence to efficient prices. We have also seen that making general statements about levels of prices and comparisons of different market structures is not as straightforward in this case and general results tend not to exist.

Comparing welfare derived in the two different informational settings, we have seen that the situation of uncertainty can be preferable when the two goods are provided by an integrated monopolist. This just aims to show consumers facing uncertainty must not necessarily be bad overall, at least when one is constrained to a world of some prices being above marginal costs.

Finally, in the appendix of the paper, I show how our main results about optimal pricing of an integrated firm depend on the correlation between primary and secondary good valuations. The integrated firm charges a secondary price above marginal costs when valuations are positively correlated, because the conditional distribution of secondary valuations is skewed to the right. The opposite holds for the case of negative correlation. The primary price tends to compensate this price change into the opposite direction.

Appendices

Appendix A: Surplus functions and lemmata

A1: Consumer surplus without uncertainty

It is the combination of both valuations that determines the purchasing decision. Since the option value is not identical for agents, an agent A with a lower primary good valuation than agent B might still demand the primary good, even if agent B does not. We will see that this will not hold when b_i is assumed to be uncertain.

Total consumer surplus derived for direct demanders in the primary market is given by:

$$CS_{1,DD}^c(p_1) = \int_{p_1}^1 (a_i - p_1)f(a_i)\,da_i.$$

In the secondary market, it is given by:

$$CS_{2,DD}^c(p_1, p_2) = (1 - F(p_1))\int_{p_2}^B (b_i - p_2)g(b_i)db_i.$$

The total surplus derived by indirect demanders can be written as:

$$CS_{ID}^c(p_1, p_2) = \int_{p_2}^B \left(\int_{p_1+p_2-b_i}^{p_1} (a_i - p_1 + b_i - p_2)f(a_i)da_i \right) g(b_i)db_i.$$

Aggregate consumer surplus in the case with certainty about b_i is the sum of all these components:

$$CS^c(p_1, p_2) = CS_{1,DD}^c(p_1) + CS_{2,DD}^c(p_1, p_2) + CS_{ID}^c(p_1, p_2).$$

Indirect demanders derive a negative surplus in the primary market, but they derive a bigger surplus as direct demanders in the secondary market, because the distribution of their types b_i is skewed towards higher b_i's. The distribution of secondary good valuations is a mix of two distributions. Direct demanders' valuations will be distributed according to $g(b_i)$ on $[0, B]$ with mass DD_1, while the distribution of types is truncated at $p_1 + p_2 - a_i$ for indirect demanders who have a mass of ID_1^c. How skewed the overall distribution is depends on the relative size of indirect demand.

A2: Producer surplus and optimal prices

Assuming constant marginal costs of c_1 form primary good production and c_2 for secondary good production, the producer surplus functions can be written as:

$$PS_1^c(p_1, p_2) = (p_1 - c_1)D_1^c(p_1, p_2)$$

$$PS_2^c(p_1, p_2) = (p_2 - c_2)D_2^c(p_1, p_2)$$

$$PS^c(p_1, p_2) = PS_1^c(p_1, p_2) + PS_2^c(p_1, p_2)$$

where $PS_1^c(p_1, p_2)$ is producer surplus in the primary market, $PS_2^c(p_1, p_2)$ is the producer surplus in the secondary market and $PS^c(p_1, p_2)$ is overall producer surplus. I am going to assume that there are no fix costs. Profit functions are, therefore, identical to producer surplus functions.

Lemma 3.

The primary producer surplus maximising primary good price is given by:

$$p_1^{PS_1^c}(p_2) = \frac{D_1^c(p_1^{PS_1^c}, p_2)}{f(p_1)G(p_2) + \int_{p_2}^B f(p_1 + p_2 - b_i)g(b_i)db_i} + c_1.$$

The secondary producer surplus maximising primary good price is:

$$p_1^{PS_2^c} = 0.$$

48

The aggregate producer surplus maximising primary good price is defined by:

$$p_1^{PS^c}(p_2) = \frac{D_1^c(p_1^{PS^c}, p_2) - (p_2 - c_2) \int_{p_2}^B f(p_1 + p_2 - b_i)g(b_i)db_i}{f(p_1)G(p_2) + \int_{p_2}^B f(p_1 + p_2 - b_i)g(b_i)db_i} + c_1.$$

The lemma implies that:

$$p_1^{PS_1^c}(p_1) > c_1$$

where the size of the mark-up depends on the primary good price and the distributional forms. Since $p_1^{PS^c}(p_2)$ takes into account externalities caused in the secondary good market, it follows that:

$$p_1^{PS^c}(p_2) \leq p_1^{PS_1^c}(p_2) \quad \text{for} \quad p_2 \geq c_2$$

$$p_1^{PS^c}(p_2) > p_1^{PS_1^c}(p_2) \quad \text{for} \quad p_2 < c_2.$$

Proof of lemma 3. The three prices defined in the lemma are the straightforward solutions to the three first order conditions below.

$$\frac{\partial PS_1^c}{\partial p_1} = D_1^c(p_1, p_2) + (p_1 - c_1) \left(\int_{p_2}^B (f(p_1) - f(p_1 + p_2 - b_i)) g(b_i)db_i - f(p_1) \right)$$

$$= D_1^c(p_1, p_2) - (p_1 - c_1) \left(f(p_1)G(p_2) + \int_{p_2}^B f(p_1 + p_2 - b_i)g(b_i)db_i \right)$$

$$\frac{\partial PS_2^c}{\partial p_1} = (p_2 - c_2) \left(\int_{p_2}^B (f(p_1) - f(p_1 + p_2 - b_i)) g(b_i)db_i - (1 - G(p_2))f(p_1) \right)$$

$$= -(p_2 - c_2) \int_{p_2}^B f(p_1 + p_2 - b_i)g(b_i)db_i$$

$$\frac{\partial PS^c}{\partial p_1} = \frac{\partial PS_1^c}{\partial p_1} + \frac{\partial PS_2^c}{\partial p_1}.$$

Since the fraction on the right hand side in the definition of $p_1^{PS_1^c}(p_2)$ in lemma 1 is greater than zero, we can say that $p_1^{PS_1^c}(p_2) > c_1$, where the extent of the mark-up depends on the own price elasticity of demand, which itself depends on the secondary price and distributional forms.

The secondary good producer surplus is decreasing in the primary good price if $\frac{\partial PS_2^c}{\partial p_1} = (p_2 - c_2)\frac{\partial D_2^c}{\partial p_1} \leq 0$, i.e. if secondary demand is decreasing in the primary good price (the goods are complements). Given that $p_2 \geq c_2$, which we verify later on, we can write this condition as:

$$\int_{p_2}^B f(p_1 + p_2 - b_i)g(b_i)db_i \geq 0$$

which is clearly always satisfied. In all cases with $f(p_1 + p_2 - b_i) > 0$ the above holds with strict

inequality for some b_i. It follows that:

$$\frac{\partial PS_2^c}{\partial p_1} \leq 0.$$

$\square$

Lemma 4.

The primary producer surplus maximising secondary good price is given by:

$$p_2^{PS_1^c} = 0.$$

The secondary producer surplus maximising secondary good price satisfies:

$$p_2^{PS_2^c}(p_1) = \frac{D_2^c(p_1, p_2^{PS_2^c})}{\int_{p_2}^B f(p_1 + p_2^{PS_2^c} - b_i)g(b_i)db_i + (1 - F(p_1))g(p_2^{PS_2^c})} + c_2.$$

The aggregate producer surplus maximising secondary good price satisfies:

$$p_2^{PS^c}(p_1) = \frac{D_2^c(p_1, p_2^{PS_2^c}) - (p_1 - c_1)\int_{p_2}^B f(p_1 + p_2 - b_i)g(b_i)db_i}{\int_{p_2}^B f(p_1 + p_2 - b_i)g(b_i)db_i + (1 - F(p_1))g(p_2)} + c_2$$

The lemma implies that:

$$p_2^{PS_2^c}(p_1) > c_2$$

where the size of the mark-up depends on the primary good price and the distributional forms. Since $p_2^{PS^c}(p_1)$ takes into account externalities caused in the primary good market, it follows that:

$$p_2^{PS^c}(p_1) \leq p_2^{PS_2^c}(p_1) \quad \text{for} \quad p_1 \geq c_1$$

$$p_2^{PS^c}(p_1) > p_2^{PS_2^c}(p_1) \quad \text{for} \quad p_1 < c_1.$$

Proof of lemma 4. Consider the effect of secondary good price changes:

$$\frac{\partial PS_1^c}{\partial p_2} = -(p_1 - c_1)\int_{p_2}^B f(p_1 + p_2 - b_i)g(b_i)db_i$$

$$\frac{\partial PS_2^c}{\partial p_2} = D_2^c(p_1, p_2) - (p_2 - c_2)\left(\int_{p_2}^B f(p_1 + p_2 - b_i)g(b_i)db_i + (1 - F(p_1))g(p_2)\right)$$

$$\frac{\partial PS^c}{\partial p_2} = \frac{\partial PS_1^c}{\partial p_2} + \frac{\partial PS_2^c}{\partial p_2}$$

The fixed point conditions defining surplus maximising secondary good prices follow directly from solving the first order conditions above. For $p_1 \geq c_1$:

$$\frac{\partial PS_1^c}{\partial p_2} = -(p_1 - c_1)\int_{p_2}^B f(p_1 + p_2 - b_i)g(b_i)db_i \leq 0$$

and vice versa, where the inequality is strict in non-trivial cases. Hence, we can say that primary good demand is decreasing in the secondary good price and, therefore, so are profits, $\frac{\partial PS_1^c}{\partial p_2} < 0$, from which the above condition about $p_2^{PS_1^c}$ in lemma 2 follows. $\qquad \square$

Appendix B: Proofs of propositions

Proof of proposition 1. When measuring welfare as the sum of consumer and producer surpluses (derived in annex A), we can derive the effects of price changes on aggregate welfare as:

$$\frac{\partial W^c}{\partial p_1} = - (p_1 - c_1)\left(f(p_1)G(p_2) + \int_{p_2}^{B} f(p_1 + p_2 - b_i)g(b_i)db_i \right)$$
$$- (p_2 - c_2)\left(\int_{p_2}^{B} f(p_1 + p_2 - b_i)g(b_i)db_i \right)$$

$$\frac{\partial W^c}{\partial p_2} = - (p_1 - c_1)\left(\int_{p_2}^{B} f(p_1 + p_2 - b_i)g(b_i)db_i \right)$$
$$- (p_2 - c_2)\left(\int_{p_2}^{B} f(p_1 + p_2 - b_i)g(b_i)db_i + (1 - F(p_1))g(p_2) \right).$$

The proposition is the unique solution to these first order conditions given our distributional assumptions (decreasing hazard rates). $\qquad \square$

Proof of proposition 2. We know that $\frac{\partial PS^c}{\partial p_1}\big|_{p_1^{PS^c}} = 0$ in optimum, in which it must also be (from the first order condition on $\frac{\partial PS^c}{\partial p_2}$ given in appendix A2, lemma 4) that:

$$-\left(1 - F(p_1^{PS^c})\right)\left((G(p_2^{PS^c})) + (p_2^{PS^c} - c_2)g(p_2^{PS^c})\right) + (p_1^{PS^c} - c_1)f(p_1^{PS^c})G(p_2^{PS^c}) = 0$$

Rewriting the above, the optimal primary price must satisfy:

$$p_1^{PS^c} = \frac{(1 - F(p_1^{PS^c}))(G(p_2^{PS^c})) + (p_2^{PS^c} - c_2)(1 - F(p_1^{PS^c}))g(p_2^{PS^c})}{f(p_1^{PS^c})G(p_2^{PS^c})} + c_1.$$

Now, suppose $p_2^{PS^c} \leq c_2$. We could then write the above condition as:

$$\overline{p}_1^{PS^c} = \frac{(1 - F(p_1^{PS^c}))}{f(p_1^{PS^c})} - \epsilon(p_2^{PS^c}, c_2) + c_1$$

where the size of ϵ depends on the level of the secondary price and the distributional forms. In any case, $\epsilon > 0$ for $p_2 < c_2$ and $\epsilon = 0$ for $p_2 = c_2$.

When comparing the above necessary condition for the primary price to the condition for $p_1^{PS_1^c}$ defined in lemma 3, we can compare the right hand side of the fixed point problems and with some

simple manipulations show that:

$$\frac{1 - F(p_1) + \int_{p_2}^{B}(F(p_1) - F(p_1 + p_2 - b_i))g(b_i)db_i}{f(p_1) - \int_{p_2}^{B}(f(p_1) - f(p_1 + p_2 - b_i))g(b_i)db_i} > \frac{1 - F(p_1)}{f(p_1)} - \epsilon \quad \forall \epsilon, p_1, p_2 \geq 0.$$

We, therefore, know that $p_1^{PS_1^c} > \overline{p}_1^{PS^c}$. Since we also know that $p_1^{PS_1^c} \leq p_1^{PS^c}$ for $p_2 \leq c_2$, we know that it must be that $p_1^{PS^c} > \overline{p}_1^{PS^c}$ under our assumption of $p_2^{PS^c} \leq c_2$, meaning that our initial assumption $(p_2^{PS^c} \leq c_2)$ cannot be valid. This proves the proposition by contradiction.

We can further say that we know that $\epsilon < 0$, $\frac{\partial \epsilon}{\partial p_2} < 0$ and $p_1^{PS_1^c} > p_1^{PS^c}$ with $\frac{\partial (p_1^{PS_1^c} - p_1^{PS^c})}{\partial p_2} > 0$ for $p_2^{PS^c} > c_2$. Hence, there exists some $p_2 > c_2$ for which $\overline{p}_1^{PS^c} = p_1^{PS^c}$.

Having established that $p_2^{PS_c} > c_2$, we can solve the second first order condition for the secondary good price and say that the profit maximising secondary price is defined by:

$$p_2^{PS^c} = \frac{(p_1^{PS^c} - c_1)f(p_1^{PS^c})G(p_2^{PS^c}) - (1 - F(p_1^{PS^c}))G(p_2^{PS^c})}{(1 - F(p_1^{PS^c}))g(p_2^{PS^c})} + c_2$$

It follows directly from the above that $p_1^{PS^c}$ is larger than the solution to $p_1 = \frac{1-F(p_1)}{f(p_1)} + c_1$, because $p_2^{PS^c} > c_2$. The optimal primary price is larger than the price charged by a monopolist in a primary market without the existence of a secondary market. $\qquad\square$

Proof of proposition 3. Suppose the integrated firm was to offer a bundle at price P. Producer surplus would then be given by:

$$(P - c_1 - c_2)\int_0^B \left(1 - F(P - b_i)\right)g(b_i)db_i.$$

The firm would derive the same producer surplus for separate pricing with $p_1 = P$ and $p_2 = 0$. However, we have shown that the optimal secondary price is given by $p_2 > c_2$. It, therefore, follows that the firm must prefer separate pricing. $\qquad\square$

Proof of proposition 5. We know that the welfare maximising prices p_1^* and p_2^* are determined by the following first order conditions:

$$\left.\frac{\partial W}{\partial p_1}\right|_{p_1=p_1^*, p_2=p_2^*} = 0$$

$$\left.\frac{\partial W}{\partial p_2}\right|_{p_1=p_1^*, p_2=p_2^*} = 0.$$

Since $\frac{\partial W}{\partial p_1}\Big|_{p_1=p_1^*, p_2=p_2^*} = 0$ in optimum, it follows from the second first order condition that:

$$-(p_2^* - c_2)(1 - F(a^{min}))g(p_2^*) = 0.$$

The second term is strictly positive in non-trivial cases and the third is positive by assumption. Therefore, it must be that the first best, welfare maximising secondary good price is given by:

$$p_2^*(c_1, c_2) = c_2.$$

This implies that the first best, welfare maximising primary good price is determined by:

$$\frac{\partial W}{\partial p_1}\Big|_{p_2=c_2} = -(p_1 - c_1)f(p_1 - EV_2(c_2)) = 0.$$

Because $f(p_1 - EV_2(c_2)) = 0$ for $p_1 < EV_2(c_2)$ and $f(p_1 - EV_2(c_2)) > 0$ otherwise, it follows that the welfare maximising price is given by:

$$p_1^*(c_1, c_2) = \begin{cases} c_1 & \text{for } c_1 > EV_2(c_2) \\ \in [0, EV_2(c_2)) & \text{for } c_1 \leq EV_2(c_2) \end{cases}.$$

$\square$

Proof of proposition 6. Rewriting the effect of a change in the secondary price, we can write the second first order conditions as:

$$\frac{\partial PS}{\partial p_2} = \frac{\partial PS}{\partial p_1}(1 - G(p_2)) - (p_2 - c_2)(1 - F(a^{min}))g(p_2).$$

Since $\frac{\partial PS}{\partial p_1}\Big|_{p_1=p_1^{PS}} = 0$ for the first first order condition to be satisfied, it follows that:

$$(p_2^{PS} - c_2)(1 - F(a^{min}))g(p_2^{PS}) = 0.$$

Since $(1 - F(a^{min})) > 0$ in non-trivial cases and $g(p_2) > 0$ by assumption, it must be that the producer surplus maximising secondary good price set by an integrated monopolist is given by:

$$p_2^{PS} = c_2.$$

Given that the secondary price is going to be equal to marginal costs, we can simplify the solution to p_1^{PS} from lemma 1 and say that p_1^{PS} is going to be defined by

$$\frac{\partial PS}{\partial p_1}\Big|_{p_1=p_1^{PS}, p_2=c_2} = 0$$

or

$$p_1^{PS} = \frac{1 - F(p_1^{PS} - E\left[V_2(c_2)\right])}{f(p_1^{PS} - E\left[V_2(c_2)\right])} + c_1.$$

Since $(1 - F(a^{min})) > 0$ for non-trivial cases and $f(a^{min}) > 0$ for $a^{min} \geq 0$ by assumption, it

follows that

$$p_1^{PS} > c_1.^{21}$$

Since $p_1^{PS} = p_1^{PS}(c_1, EV_2(p_2))$, we can differentiate the above solution to show that:

$$\frac{\partial p_1^{PS}}{\partial c_1} = \frac{1}{2 + (1 - F(\cdot))\frac{f'(\cdot)}{f(\cdot)^2}}$$

$$\frac{\partial p_1^{PS}}{\partial EV_2} = 1 + \frac{(1 - F(\cdot))f'(\cdot)}{f(\cdot)^2}.$$

Under our assumption of an increasing hazard rate function, it follows directly that:

$$1 > \frac{\partial p_1^{PS}}{\partial c_1} > 0$$

$$\frac{\partial p_1^{PS}}{\partial EV_2} > 0.$$

$\square$

Proof of proposition 7. Suppose the monopolist charges a price P for the bundle. All consumers with $a_i + EV_2(0) \geq P$ would purchase the bundle. Producer surplus would then be given by:

$$PS(P) = (P - c_1 - c_2)\big(1 - F(a_i + EV_2(0))\big).$$

Now, suppose that the monopolist uses separate pricing and charges a primary good price of $p_1 = P$ and a secondary good price of $p_2 = 0$. Producer surplus would be the same as above. However, we have shown that $p_2^{PS} = c_2$. It follows directly that:

$$PS(p_1^{PS}, c_2) \geq PS(P)$$

where the only time the above holds with equality is for $c_2 = 0$.

$\square$

---

[21]This result is similar to the case of a monopolist using a two-part tariff to extract consumer surplus when consumers are homogeneous. In that scenario, a monopolist charges a per unit price equal to marginal costs and extracts consumer surplus by charging a high lump-sum fee. This result resembles that pricing strategy, because, although consumers are not homogeneous themselves, their option values are identical. However, although the result resembles a two-part tariff with homogeneous agents, it is derived in an entirely different setting. Consumers are not homogeneous and there are usually no asymmetric complementarities in models using a two-part tariff. Although the membership fee is necessary for further consumption, it is typically not a stand alone product providing utility on its own. In addition, both goods in our model are consumed in fixed proportions.

Proof of proposition 9. An equilibrium is a set of prices $p_1^{PS_1}$ and $p_2^{PS_2}$ such that:

$$p_1^{PS_1} = \frac{1 - F(a^{min})}{f(a^{min})} + c_1$$

$$p_2^{PS_2} = \frac{(1 - F(a^{min}))(1 - G(p_2^{PS_2}))}{(1 - F(a^{min}))g(p_2^{PS_2}) + f(a^{min})(1 - G(p_2^{PS_2}))^2} + c_2.$$

Hence, in equilibrium the secondary good price is strictly larger than marginal costs and, therefore, higher than in the integrated firm case. The primary good price is lower than in the integrated firm case, because the primary reaction function is decreasing in the secondary price. The absolute and relative sizes of the differences depend on the distribution functions and the relative value of the secondary good.

However, we can say that the sum of prices under competition is higher, because $\frac{\partial p_1^{PS_1}}{\partial p_2} \geq -1$. This, on its own, does not tell us in which case consumer surplus is higher, because higher secondary prices are also less harmful to consumer surplus. However, we it is clear that the primary good price decreases by less than necessary in order to compensate for the decrease in option value, $\frac{\partial p_1^{PS_1}}{\partial p_2} \geq -(1 - G(p_2))$, which means that consumer surplus is smaller under competition than under an integrated firm. It follows that welfare must also be lower under Bertrand competition, because producer surplus must have gone down. The integrated firm could have chosen the prices derived under Bertrand-like competition, but chose not to do so. Producer surplus must, therefore, be lower under competition across markets. □

Proof of proposition 10. We can see how the consumer surplus depends on the value of the secondary good:

$$\frac{dCS}{dB} = (1 - F(a^{min})) \left(\frac{\partial EV_2}{\partial B} + \frac{\partial EV_2}{\partial p_2} \frac{\partial p_2}{\partial B} - \frac{\partial p_1}{\partial B} \right).$$

In the absence of price changes induced by a change in B (disregarding the firms response to a change in B) it can be written:

$$\frac{\partial CS}{\partial B} = (1 - F(a^{min})) \frac{\partial EV_2}{\partial B}.$$

In this case, the entire increase in the option value goes to consumers.

In general, one has to take into consideration that firms adjust their prices. Part of the additional surplus created by an increase in B will, therefore, be absorbed by producers. Differentiating the implicit solution for the optimal primary price and using the fact that the optimal secondary good price of the integrated firm is always equal to marginal costs ($p_2^{PS} = c_2$), which implies that $\frac{\partial p_2^{PS}}{\partial B} = 0$, we can express how the optimal primary good price charged by the integrated firm

depends on B:

$$\frac{\partial p_1^{PS}}{\partial B} = \frac{(1 - F(a^{min}))f'(a^{min}) + f(a^{min})^2}{(1 - F(a^{min}))f'(a^{min}) + 2f(a^{min})^2} \frac{\partial EV_2}{\partial B} < \frac{\partial EV_2}{\partial B}.$$

Plugging this into the expression for the change in consumer surplus, we can say that consumers surplus, once price changes are taken into account, reacts as follows for the case of an integrated firm:

$$\frac{\partial CS^I}{\partial B} = (1 - F(a^{min}))\frac{f(a^{min})^2}{(1 - F(a^{min}))f'(a^{min}) + 2f(a^{min})^2} \frac{\partial EV_2}{\partial B}.$$

We can also calculate how the producer surplus reacts to changes in the secondary good value:

$$\frac{\partial PS}{\partial B} = (1 - F(a^{min})) \left(\frac{\partial p_1}{\partial B} + (1 - G(p_2))\frac{\partial p_2}{\partial B} - (p_2 - c_2)\left(\frac{\partial G(p_2)}{\partial B} + \frac{\partial G(p_2)}{\partial p_2}\frac{\partial p_2}{\partial B} \right) \right)$$
$$- \left((p_1 - c_1) + (1 - G(p_2))(p_2 - c_2) \right) \left(f(a^{min})\left(\frac{\partial p_1}{\partial B} - \frac{\partial EV_2}{\partial B} - \frac{\partial EV_2}{\partial p_2}\frac{\partial p_2}{\partial B} \right) \right).$$

Hence, if the firm/firms were not adjusting their prices, the change in producer surplus would be given by:

$$\frac{\partial PS}{\partial B} = -(p_2 - c_2)(1 - F(\cdot))\frac{\partial G(p_2)}{\partial B} + \left((p_1 - c_1) + (1 - G(p_2))(p_2 - c_2) \right)f(\cdot)\frac{\partial EV_2}{\partial B} > 0.$$

Producer surplus would still increase, because demand increases due to an increase in the option value $\left(\frac{\partial G(p_2)}{\partial B} < 0 \right)$.

We know that $\frac{\partial p_2^{PS}}{\partial B} = 0$ and we can, therefore, say that:

$$\frac{\partial PS^I}{\partial B} = (1 - F(a^{min}))\frac{\partial p_1^{PS}}{\partial B} - (p_1 - c_1)f(a^{min})\left(\frac{\partial p_1^{PS}}{\partial B} - \frac{\partial EV_2}{\partial B} \right),$$

This expression is strictly positive as can be seen by the fact that $\frac{\partial EV_2}{\partial B} > \frac{\partial p_1^{PS}}{\partial B} > 0$ or by plugging in the specific solution for $\frac{\partial p_1^{PS}}{\partial B}$ derived above.

Having derived the effects on consumer and producer surplus functions, we can combine these and derive the effect on aggregate welfare, which follows immediately from the above as our measure of welfare is the pure sum of consumer and producer surplus.

The change in aggregate welfare in the case of an integrated firm, taking into account price reactions by the firm, is given by:

$$\frac{\partial W^I}{\partial B} = (1 - F(a^{min}))\frac{\partial EV_2}{\partial B} - (p_1 - c_1)f(a^{min})\left(\frac{\partial p_1}{\partial B} - \frac{\partial EV_2}{\partial B} \right).$$

In the final step of the proof, we want to compare this to the behaviour of first best welfare. We have seen that $p_1 = c_1$ and $p_2 = c_2$ always, independent of B, maximise total welfare. The first best prices do not respond to the value of the secondary good and the behaviour of first best surpluses

can, therefore, be summarised as follows:

$$\frac{\partial CS^*}{\partial B} = (1 - F(a^{min}))\frac{\partial EV_2}{\partial B}$$
$$\frac{\partial PS^*}{\partial B} = 0$$
$$\frac{\partial W^*}{\partial B} = (1 - F(a^{min}))\frac{\partial EV_2}{\partial B}.$$

Furthermore, we know that the integrated firm would act as a simple monopolist for $B = 0$, so that:

$$CS^I|_{B=0} < CS^*|_{B=0}$$

$$PS^I|_{B=0} > PS^*|_{B=0}$$

$$W^I|_{B=0} < W^*|_{B=0}.$$

Compare now $\frac{\partial W^*}{\partial B}$ with $\frac{\partial W^I}{\partial B}$:

$$\frac{\partial W^I}{\partial B} - \frac{\partial W^*}{\partial B} = -(p_1 - c_1)f(a^{min})\left(\frac{\partial p_1}{\partial B} - \frac{\partial EV_2}{\partial B}\right)$$

This expression is strictly positive if $p_1 > c_1$ and $\frac{\partial p_1^{PS}}{\partial B} < \frac{\partial EV_2}{\partial B}$, which we have shown to be the case. The the optimal primary good price charged by an integrated firm increases with the secondary good value, but by less than the option value itself. The entire expression becomes zero once we are in full demand, because $f(a^{min}) = 0$ for $a^{min} < 0$. Welfare is, therefore, increasing faster in the integrated case as long as not in full demand:

$$\frac{\partial W^I}{\partial B} > \frac{\partial W^*}{\partial B} \quad \forall p_1 > EV_2$$

$$\frac{\partial W^I}{\partial B} = \frac{\partial W^*}{\partial B} \quad \forall p_1 \leq EV_2.$$

These results are summarised in the proposition. $\square$

Appendix C: Extension - Correlated valuations

Another assumption that has potentially been driving results is that valuations for primary and secondary goods were assumed to be uncorrelated. I will consider the case of correlated valuations mainly for the case of uncertain secondary valuations. In the case without uncertainty, the case is a bit trivial, because correlation does not change the individual decision process and all changes to the allocation are due to distributional changes.[22]

---

[22]Furthermore, most of the interesting results were derived in the setting with uncertainty and it seems more interesting to analyse how these results stand up to correlated valuations.

No uncertainty

The distribution of secondary valuations depends on the primary valuations, $G_{a_i}(p_2)$. We can, thus, rewrite aggregate demands as when faced with a correlation of ρ as:

$$DD_{1,\rho}(p_1) = 1 - F(p_1) = DD_1(p_1)$$

$$ID_{1,\rho}(p_1, p_2) = \int_{p_2}^{B} \left(\int_{p_1+p_2-b_i}^{p_1} f(a_i)da_i \right) g_{a_i}(b_i)db_i$$

$$D_{2,\rho}(p_1, p_2) = ID_{1,\rho}(p_1, p_2) + \int_{p_1}^{1} (1 - G_{a_i}(p_2))f(a_i)da_i$$

and proceed as earlier. It is worth mentioning that indirect demand is obviously larger the more negative the correlation between the valuations.

Uncertainty

The main implication of correlated valuations with uncertainty is that the option value of the primary good is not identical for all consumers. It depends on the primary valuation a_i and the level of correlation via the distribution function for secondary good valuations.[23]

Consumers will demand the primary good if $a_i + EV_2^{\rho}(p_2) \geq p_1$, with:

$$\frac{\partial EV_2^{\rho}}{\partial a_i} < 0 \text{ for } \rho < 0$$

$$\frac{\partial EV_2^{\rho}}{\partial a_i} > 0 \text{ for } \rho > 0.$$

Dropping the assumption of no correlation (as was the case with dropping uncertainty) does not affect the decision of direct demanders. Direct demand under correlation DD_1^{ρ} is unchanged and given by:

$$DD_1^{\rho}(p_1) = (1 - F(p_1)) = DD_1(p_1)$$

Assume, for simplification, that the correlation is never very negative. More specifically, assume that $\frac{\partial a_i + EV_2^{\rho}(\cdot)}{\partial a_i} \geq 0$ or $\frac{\partial EV_2^{\rho}(\cdot)}{\partial a_i} \geq -1$. Otherwise, the agents' total willingness to pay would not be monotonically increasing in the primary good valuation. If this assumption was violated, which could be the case if B was large and the correlation was very negative at the same time, the set of consumers demanding the primary good would not necessarily be convex in the type set a_i. If $\frac{\partial a_i + EV_2^{\rho}(\cdot)}{\partial a_i} < 0$ for all a_i, then there would exist a threshold a^{max} for which $a^{max} + EV_2^{\rho}(\cdot) = p_1$ and below which everyone demands the primary good, because the willingness to pay would be decreasing rather than increasing in a_i. If, for example, $\frac{\partial a_i + EV_2^{\rho}(\cdot)}{\partial a_i} < 0$ for low values of a_i and vice versa for large a_i, then there would exist a threshold a^{max} below which everyone is an indirect

---

[23]Note that we showed that the results derived under uncertainty depended on the fact that the option value was constant. Hence, we can already make an educated guess that the results derived hold only for zero correlation.

demander. Everyone above $a_i = p_1$ would be a direct demander.

The cases of positive and negative correlation and their effect on consumer demand are shown in figure 17. It includes the two cases of non-monotonically increasing willingness to pay in red. The figure shows how the decreasing willingness to pay with very negative correlation and high B leads to a threshold a^{max} rather than a^{min} and possibly to a non-convex set of demanders on the type space a_i.

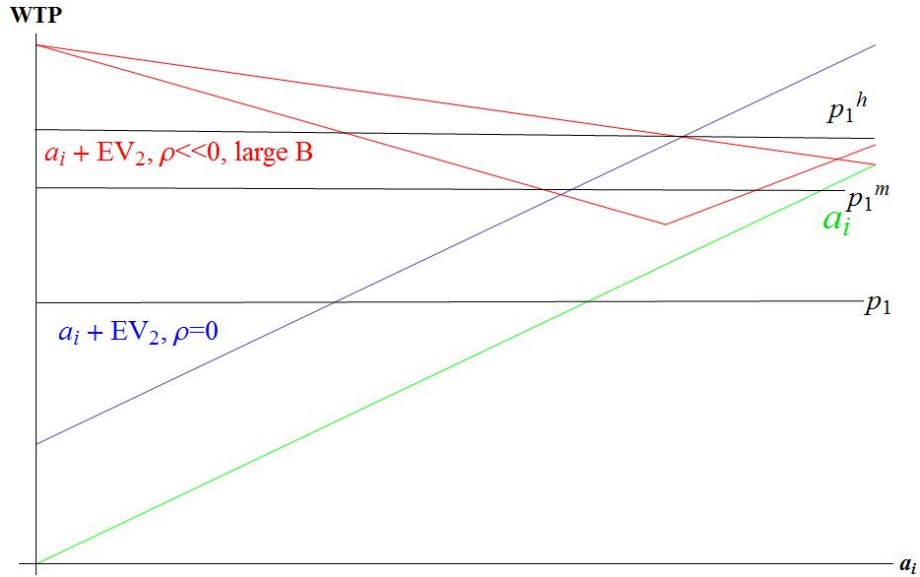Figure 17: Potentially non-convex set of demand

In figure 18, depicting the three standard cases we allow for under our assumptions, we can see that direct demand is independent of the level of correlation. Furthermore, we can see that at relatively low primary good prices demand is increasing in ρ while for high prices the opposite is the case. The slope of the willingness to pay is increasing in ρ.
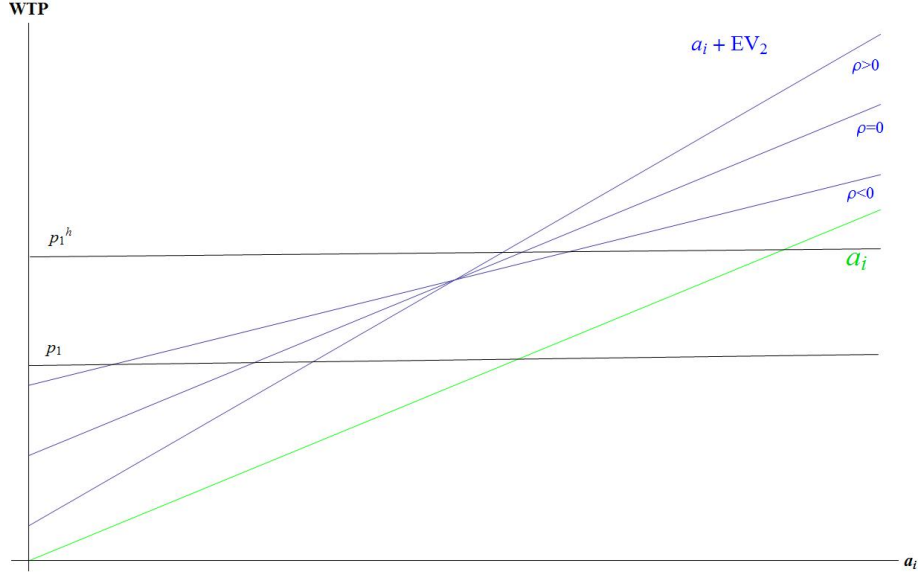
Figure 18: The effect of correlation on willingness to pay

With $EV_2^\rho = \int_{p_2}^{B}(b_i - p_2)g_{a_i}(b_i)db_i$ and our monotonicity assumptions about $a_i + EV_2^\rho(\cdot)$, we can be ensured that indirect demand is given by:

$$ID_1^\rho(p_1, p_2) = F(p_1) - F(a^{min,\rho}),$$

where $a^{min,\rho} = p_1 - EV_2(p_2)$ is defined uniquely.

We can, therefore, rewrite demand functions and producer surplus functions in a more general setting, allowing for correlation:

$$D_1^\rho(p_1, p_2) = (1 - F(a^{min,\rho}))$$
$$D_2^\rho(p_1, p_2) = \int_{a^{min,\rho}}^{1}(1 - G_{a_i}(\cdot))f(a_i)da_i$$
$$PS_1^\rho(p_1, p_2) = (p_1 - c_1)(1 - F(a^{min,\rho}))$$
$$PS_2^\rho(p_1, p_2) = (p_2 - c_2)\int_{a^{min,\rho}}^{1}(1 - G_{a_i}(\cdot))f(a_i)da_i.$$

Hence, we can write the marginal profits for the integrated firm as:

$$\frac{\partial PS^\rho(\cdot)}{\partial p_1} = (1 - F(a^{min,\rho})) - (p_1 - c_1)f(a^{min,\rho})\frac{\partial a^{min,\rho}}{\partial p_1}$$
$$- (p_2 - c_2)(1 - G_{a^{min,\rho}}(\cdot))f(a^{min,\rho})\frac{\partial a^{min,\rho}}{\partial p_1}$$

$$\frac{\partial PS^\rho(\cdot)}{\partial p_2} = \int_{a^{min,\rho}}^{1}(1 - G_{a_i}(\cdot))f(a_i)da_i - (p_1 - c_1)f(a^{min,\rho})\frac{\partial a^{min,\rho}}{\partial p_2}$$
$$- (p_2 - c_2)\left((1 - G_{a^{min,\rho}}(\cdot))f(a^{min,\rho})\frac{\partial a^{min,\rho}}{\partial p_2} + \int_{a^{min,\rho}}^{1}g_{a_i}(\cdot)f(a_i)da_i\right)$$

Differentiating both sides of the definition of $a^{min,\rho}$ we get:

$$\frac{\partial a^{min,\rho}}{\partial p_1} = \frac{1}{1 + \int_{p_2}^{B}(b_i - p_2)\frac{\partial g_{a^{min,\rho}}(\cdot)}{\partial a^{min,\rho}}db_i}$$

$$\frac{\partial a^{min,\rho}}{\partial p_2} = (1 - G_{a^{min,\rho}}(\cdot))\frac{\partial a^{min,\rho}}{\partial p_1}.$$

This shows that a change in the secondary good price changes the threshold valuation by a smaller amount than primary price changes, because it only works via a change in the option value.

We know that $\frac{\partial g_{a^{min}}(\cdot)}{\partial a^{min}} = 0$ and, hence, $\frac{\partial a^{min}}{\partial p_1} = 1$ in the case of no correlation. As shown earlier, the marginal effects caused by a secondary price change are just a fraction of the ones caused by primary price changes.

Rewriting the marginal profits in the same way as before, we get:

$$\frac{\partial PS^\rho}{\partial p_2} = \frac{\partial PS^\rho}{\partial p_1}(1 - G_{a^{min,\rho}}(\cdot))$$

$$+ \left[\int_{a^{min,\rho}}^{1}(1 - G_{a_i}(\cdot))f(a_i)da_i - (1 - G_{a^{min,\rho}})(1 - F(a^{min,\rho}))\right]$$

$$- (p_2 - c_2)\int_{a^{min,\rho}}^{1} g_{a_i}(p_2)f(a_i)da_i.$$

Remember how the distribution of secondary valuations conditional on primary good demand in the case of certainty was skewed towards higher valuations. This caused a secondary good price above marginal costs. The same logic holds for the case of non-zero correlation. For positive correlation, $\rho > 0$, the term in the square brackets above, which shows the additional benefit caused by the rightward skewness of secondary valuations, is positive, resulting in a secondary good price above marginal costs. The opposite holds for negative correlation, resulting in a secondary good price below marginal costs.

In general, rewriting the above, knowing that $\frac{\partial PS^\rho}{\partial p_1}|_{p_1=p_1^{PS,\rho},p_2=p_2^{PS,\rho}} = 0$, we can say that the optimal secondary good price, $p_2^{PS,\rho}$ is given by:

$$p_2^{PS,\rho} = \frac{\int_{a^{min,\rho}}^{1}(1 - G_{a_i}(\cdot))f(a_i)da_i - (1 - G_{a^{min,\rho}})(1 - F(a^{min,\rho}))}{\int_{a^{min,\rho}}^{1} g_{a_i}(p_2)f(a_i)da_i} + c_2.$$

This implies the following result for the optimal secondary price charged by an integrated firm.

The producer surplus maximising secondary good price and, thus, the price chosen by an integrated firm, when there is uncertainty about secondary good valuations, is given by:

$$p_2^{PS,\rho} \begin{cases} < c_2 & \text{for } \rho < 0 \\ = c_2 & \text{for } \rho = 0 \\ > c_2 & \text{for } \rho > 0 \end{cases}.$$

Rewriting and simplifying the primary price marginal profit condition, we can say that the optimal primary good price, $p_1^{PS,\rho}$, is given by:

$$p_1^{PS,\rho} = \left(1 + \int_{p_2}^{B}(b_i - p_2)\frac{\partial g_{a^{min,\rho}}(b_i,\cdot)}{\partial a^{min,\rho}}db_i\right)\frac{1 - F(\cdot)}{f(\cdot)} - (p_2^{PS,\rho} - c_2)(1 - G_{a^{min,\rho}}(\cdot)) + c_1.$$

Combining the optimal pricing rule, $p_2^{PS,\rho}$ with the fact that the first term in the brackets, which is just the inverse of $\frac{\partial a^{min}}{\partial p_1}$, is increasing in ρ, the two effects affect the optimal primary good price in opposite ways. Whether the primary good price increases or decreases with correlation, therefore, depends on the specified distributions and parameter values.

To improve intuition, we can consider deviations from existing prices like we did for the intuitive proof in the case of zero correlation. Consider a firm that currently charges prices p_1' and p_2' and let us consider a change of prices that leaves demand in the primary market unchanged. In other words, consider a change of prices that leaves $a^{min,\rho}$ unchanged. The ratio of primary to secondary price which achieves this, as shown before, is given by $G_{a^{min,\rho}}(\cdot) - 1$.

Consider a change of the secondary good price of ϵ and the corresponding primary price change $(G_{a^{min,\rho}}(\cdot) - 1)\epsilon$, and evaluate the change in producer surplus the producer derives from consumer type a_i, ΔPS_{a_i}:

$$\Delta PS_{a_i} = \Delta(G_{a^{min,\rho}}(\cdot) - 1)\epsilon D_1(a_i) - (p_2' + \epsilon - c_2)g_{a_i}(\cdot)\epsilon D_1(a_i) + (1 - G_{a^{min,\rho}}(\cdot))\epsilon D_1(a_i).$$

Simplifying further and noticing that $D_1 = 0$ for $a_i < a^{min,\rho}$ and $D_1 = 1$ otherwise, we can say that $\Delta PS_{a_i} = 0$ for $a_i < a^{min,\rho}$ and:

$$\Delta PS_{a_i \geq a^{min,\rho}} = \epsilon\left[G_{a^{min,\rho}}(\cdot) - G_{a_i}(\cdot) - (p_2' + \epsilon - c_2)g_{a_i}(\cdot)\right].$$

In the case of no correlation $G_{a^{min,\rho}}(\cdot) = G_{a_i}(\cdot) = G(\cdot)$ and, thus, $\Delta PS_{a_i \geq a^{min,\rho}} = \Delta PS_{a_j \geq a^{min,\rho}} = \Delta PS$ and we get our earlier result:

$$\Delta PS = -\epsilon(p_2' + \epsilon - c_2)g(\cdot).$$

Hence, whenever the secondary price is above marginal costs a price decrease increases profits derived from each type a_i and, therefore, overall producer surplus. The reverse is true when the secondary price is below marginal costs.

Now consider valuations that are positively correlated, $G_{a^{min,\rho}} > G_{a_i > a^{min,\rho}}$. Consumers with high primary valuations are less likely to have small secondary good valuations. We can write:

$$\Delta PS_{a^{min,\rho}} = \epsilon\left[G_{a^{min,\rho}}(\cdot) - G_{a^{min}}(\cdot) - (p_2' + \epsilon - c_2)g_{a^{min,\rho}}(\cdot)\right]$$

$$= -\epsilon(p_2' + \epsilon - c_2)g_{a^{min,\rho}}(\cdot)$$

$$\Delta PS_{a_i} = \epsilon \left[G_{a^{min,\rho}}(\cdot) - G_{a_i}(\cdot) - (p_2' + \epsilon - c_2) g_{a_i}(\cdot) \right]$$

$$= -\epsilon((p_2' + \epsilon - c_2) g_{a_i}(\cdot) + G_{a_i}(\cdot) - G_{a^{min,\rho}}(\cdot)).$$

The total change in producer surplus can then be written as:

$$\Delta PS = \int_{a^{min,\rho}}^{1} -\epsilon((p_2' + \epsilon - c_2) g_{a_i}(\cdot) + G_{a_i}(\cdot) - G_{a^{min,\rho}}(\cdot)) da_i$$

Consider the case where the firm initially charges marginal costs in the secondary market, which was optimal with zero correlation. In that case, we can write the above as:

$$\Delta PS = \int_{a^{min,\rho}}^{1} -\epsilon(\epsilon g_{a_i}(\cdot) - \xi(a_i, a^{min,\rho})) da_i$$

where $\xi(a_i, a^{min,\rho}) = G_{a^{min,\rho}}(\cdot) - G_{a_i}(\cdot)$ is positive and increasing in a_i, precisely because higher primary valuations imply higher expected secondary valuation.

$\xi(\cdot) = 0$ in the case of no correlation. This meant that no profitable deviation existed. However, in the case of positive correlation, we can clearly see that the term in the brackets is smaller than zero for all $a_i > a^{min,\rho}$ and zero for $a^{min,\rho}$ for a sufficiently small $\epsilon > 0$. An increase in the secondary good price is profit increasing and the optimal secondary good price must, therefore, be above marginal costs. Whatever a^{min} is, which is determined by the optimal prices, the ratio of prices must be such that the equation above is zero, which is at some $p_2 > c_2$.

Intuitively, an increase in p_2 causes some demanders with lower a_i to stop demanding the secondary good, but this negative demand effect is smaller for consumers with high a_i because of the positive correlation. The more positive the correlation, the larger $\xi(a_i, a^{min,\rho})$ for all a_i and the larger the price mark-up over marginal costs.

In the case of negative correlation the exact opposite is true $(G_{a^{min,\rho}} < G_{a_i > a^{min,\rho}})$. The change in producer surplus just changes its sign and can be written as:

$$\Delta PS = \int_{a^{min}}^{1} -\epsilon(\epsilon g_{a_i}(\cdot) - \xi(a_i, a^{min})) da_i.$$

A deviation involving a secondary good price decrease is beneficial, because $\xi(\cdot)$ is negative for $\rho < 0$ and decreasing in a_i.

Total valuations differ very little across consumers ex post in the case of negative correlation. In response, the integrated firm creates a high option value for consumers, which enables the firm to charge a high primary good price. Very few of these consumers are going to have a secondary valuation ex post such that they will actually demand the secondary good. The downside from charging below marginal costs in the secondary market is, therefore, low compared to the upside of the increased option value. In the case of positive correlation, the total valuations differ extremely across consumers. Decreasing the option value by charging a high secondary good price comes at

little cost, because most consumers moving into the secondary market will demand the secondary good.

References

Adams, William J.; and Yellen, Janet L. "Commodity Bundling and the Burden of Monopoly", *The Quarterly Journal of Economics*, 1976, 90(3), pp. 475-498

Andriychenko, Oleksiy; Girnius, Audrius; and Saha, Atanu. "Complementary Goods: Prices and Consumer Welfare Under Duopoly and Monopoly", *international Journal of the Economics of Business*, 2006, 13(3), pp. 373-386

Blair, Roger D.; and Kaserman, David L. "Vertical Integration, Tying, and Antitrust Policy", *The American Economic Review*, 1978, 68(3), pp. 397-402

Cheng, Leonard K.; and Nahm, Jae. "Product Boundary, Vertical Competition and the Double Mark-Up Problem", *The RAND Journal of Economics*, 2007, 38(2), pp. 447-466

Cheng, Leonard K.; and Nahm, Jae. "Asymmetric Complementary Goods Pricing under Sequential Moves", *The B.E. Journal of Economic Analysis & Policy*, 2010, 10(1), Article 46

Davis, Steven J.; and Murphy, Kevin M. "A competitive perspective on Internet Explorer", *The American Economic Review: P&P*, 2000, 90(2), pp. 184-187

Economides, Nicholas; and Salop, Steven C. "Competition and Integration Among Complements, and Network Market Structure", *The Journal of Industrial Economics*, 1992, 40(1), pp. 105-123

Economides, Nicholas; and Viard, V. Brian. "Pricing of Complements and Network Effects", 2009

Ellison, Glenn. "A Model of Add-On Pricing", *The Quarterly Journal of Economics*, 2005, 120(2), pp. 585-637

Farrell, Joseph; and Katz, Michael L. "Innovation, Rent Extraction, and Integration in Systems Markets", *The Journal of Industrial Economics*, 2000, 48(4), pp. 413-432

Gans, Joshua S. "Remedies for tying in computer applications", *International Journal of Industrial Organization*, 2011, 29, pp. 505-512

Lerner, Josh; and Tirole, Jean. "Efficient Patent Pools", *The American Economic Review*, 2004, 94(3), pp. 691-711

Matutes, Carmen; and Regibeau, Pierre. "Mix and Match: Product Compatibility without Network Externalities", *The RAND Journal of Economics*, 1999, 19(2), pp. 221-234

Matutes, Carmen; and Regibeau, Pierre. "Compatibility and Bundling of Complementary Goods in a Duopoly", *The Journal of Industrial Economics*, 1992, 40(1), pp. 37-54

McAfee, Preston R.; McMillan, John; and Whinston, Michael D. "Multiproduct Monopoly, Commodity Bundling, and Correlation of Values", *The Quarterly Journal of Economics*, 1989, 104(2), pp. 371-383

Rey, Patrick; and Tirole, Jean. "Cooperation vs. Collusion: How Essentiality Shapes Coopetition", *IDEI*, 2013, WP 801

Whinston, Michael D. "Tying, Foreclosure, and Exclusion", *The American Economic Review*, 1990, 80(4), pp. 837-859

The Effect of Competition and Dispersion of Public Opinion on Media Bias Over Time

Kimon Doulis
School of Economics
University of Edinburgh

Abstract

The paper presents a model in which media and interest groups seek to maximise their support in the population, but also aim to further their own agenda as much as possible by providing potentially biased reports. Reports are not sold to agents at a price, but are freely available to followers, who have to decide which group to follow. The private views of these followers change over time, depending on the news reports they received.

A monopolistic media group, when faced with a homogeneous group of followers, will, in equilibrium, bias its news. It will assign some weight to its agenda, some to the followers' priors and some to to the truth as long as slanting is costly. Introducing competition eliminates the motive to bias information according to the group's own agenda. It is unclear whether overall bias is decreasing in the level of competition; this depends on the relative extremity of media groups' agendas and the priors of followers. However, bias decreases over time under competition, while its evolution in the monopolistic setting is ambiguous. This means that competition becomes more favourable over time. Public opinion converges to the true state of the world under competition. This is not the case in the monopolistic setting.

Optimal bias provided by a monopolistic group when faced with heterogeneous followers is similar to the homogeneous case, with slight differences depending on the distribution of priors. Competition does not eliminate the agenda motive of media groups when the groups are faced with heterogeneous followers. Overall, competition either has no effect in terms of optimal slanting policies or creates an incentive to increase bias. This incentive to increase bias is potentially worse when confirmation bias is increasing in the extremity of priors. Competition can, therefore, increase dispersion of public opinion, and there is no real sense of competition becoming more preferable over time when followers are heterogeneous.

1 Introduction & Literature

Members of society have different points of view on a variety of issues. This can very much enrich public debate. However, there is arguably some level of dispersion that is harmful, as it will become increasingly hard to create compromise between different sections of society. When there is some underlying factual truth, there is value to members of society agreeing upon this truth. A society

will find this increasingly hard if it becomes highly segmented. Public discourse and decision-making can worsen if the dispersion of private opinions is large. In addition, individual opinions and public opinion are not constants. They develop over time. Their development depends largely on what information people receive.[1] One important source of information, based on which people might adapt their point of view, are reports by media and interest groups.

Supply side models

When analysing the behaviour of interest groups, the traditional approach has been to model them as having a specific own agenda. The groups seek to find support for their agenda in the population and try to influence public decision making by either contributing monetarily to pivotal policy makers (monetary lobbying) or providing biased information to them (informational lobbying). Bennedsen and Feldmann (2005) present a model highlighting the incentive for lobby groups to provide biased information to policy makers under different political systems. There might also be a relationship between the lobby group providing biased information to policy makers and contributions it makes to gain access to policy makers (Austen-Smith, 1995).

The reason for equilibrium bias in these lobbying models is essentially the same as in models of supply side driven media bias. Biases arise because the private interests of lobbies (or media groups) are not aligned with the interests of policy makers (or the public in general). In these supply side models of media bias, biases arise from preferences of owners and stakeholders (Anderson and McLaren, 2010) or journalists (Baron, 2006). Readers are assumed to be rational in these models. They do not suffer from any form of self-serving bias and update prior beliefs in a Bayesian way. They take potential media bias into account when updating beliefs. Information provided by profit maximising media groups can be biased in equilibrium, even though readers are rational, because of information asymmetries. Readers simply do not know how much information the media groups have and whether they are withholding disadvantageous information just as policy makers do not know what information lobby groups are withholding in Bennedsen and Feldmann (2005).

Demand side models

A different part of the literature on media bias can be described as demand side models. Equilibrium bias can be sustained in these models typically because of confirmation bias, which means that readers have an intrinsic preference for their prior views to be confirmed. Mullainathan and Shleifer (2005) and Bernhardt et al. (2008) show that media groups tend to cater towards readers' priors

---

[1] It is not straightforward to define the term 'public opinion'. In this paper 'public opinion' is not defined as a basic measure such as the average private opinion, or even the opinion with the largest support. The dispersion of private opinions matter and public opinion is, therefore, described by the entire distribution of individual opinions.

when readers dislike hearing news that differ from their views. Bernhardt et al. (2008), in a static model, show how media bias can lead to an increased probability of electoral mistakes because of increased political polarisation. Models of media groups or politicians 'pandering' to individuals' priors and confirmation biases fall into this category of demand side models of bias. Ashworth and Shotts (2010) present such a paper in which politicians have an incentive to implement bad policies against their better knowledge if voters believe such a policy to be best for them. This incentive can be offset or increased by the existence of media groups, depending on the motives of such a media group. Media groups could have an incentive to pander themselves or side with the popular politician to increase own popularity. The incentive for politicians to pander to voters when seeking re-election has also been highlighted by Maskin and Tirole (2004). Finally, Gentzkow and Shapiro (2006) show that such demand driven media bias can arise even when readers do not suffer from confirmation bias and are Bayesian. When media groups differ in terms of quality and readers cannot observe quality, media reports tend to cater to priors of readers, because they are, on average, seen as being more likely to be of high quality.

Most of the recent literature (in particular Bernhardt et al, 2008; Gentzkow and Shapiro, 2006; Anderson and McLaren, 2010) focuses on rational, Bayesian readers and takes the point of view that results derived in Mullainathan and Shleifer (2005) depend on irrational readers and could, thus, be undone by introducing rational readers. However, the existence of preferences for confirmatory news seems well established. Assuming that it does not exist, seems to turn the analysis into a technical exercise that ignores the existence of such biases. Both economic and psychology literatures show that agents suffer from confirmation bias in the selection of their sources (Jones and Sugden, 2001; Severin and Tankard, 1992; Nickerson, 1998) as well as in evaluating information (Lord et al, 1979; Zaller, 1992; Rabin, 1999; Mynatt et al, 1977). These papers show that agents tend to seek information that confirms their views and to disregard opposing views to an extent. They also suggest that information is used for updating mainly when it is confirming pre-existing views.[2] Furthermore, the literature (especially Lord et al, 1979 and Zaller, 1992) hints at the fact that this confirmatory bias is stronger, the stronger the agents viewpoint. In other words, people who really care and have strong views about a specific issue tend to suffer from a stronger confirmation bias than people with moderate views on the issue.

Although there seems to be some evidence for Bayesian updating as well that indicates that readers do not just blindly update their beliefs towards the news source they are presented with (Gerber et

---

[2]A recent survey by the Pew Research Center, conducted on 20-24 July 2011 shows similar results. Only 25% of respondents said that news organisations get the facts straight, while the overall majority (66%) think that news stories are often inaccurate. When asked the same question about the news outlet they mostly used, opinions shifted with 62% claiming that it gets the facts straight and only 30% admitting that stories are often inaccurate.

al, 2009), the evidence is more limited.[3] In this paper, I do not assume Bayesian readers. Instead, readers (which I refer to as followers) change their views according to a constant parameter. If wished, one could endogenise this parameter. I do not attempt this here for simplification and tractability issues. The evidence supporting the idea that readers are Bayesian is also, as pointed out, limited.

Combining demand and supply side models

It is fairly established that people like to hear news that confirm their beliefs. This has been used to explain why media groups might want to provide biased information and slant their reports towards readers' priors in the demand side literature. It remains to be asked how groups in general, whether media or interest groups, trade off confirmation bias of potential followers and own agenda.

I will combine the supply and demand side motives for bias into a single model and try to explain when which motives dominates. Gentzkow and Shapiro (2010) using US data show that media outlets respond strongly to readers' priors and that the demand side motive is, thus, much stronger at explaining bias than the supply side motive. The model presented here will provide a potential theoretical explanation for these findings by arguing that the supply side motive is weakened much more than the demand side motive under competition.

The relationship between media and interest groups

The traditional approach in the informational lobbying literature, as explained above, has been that these groups collect beneficial information from the wider public and provide it to policy makers in a biased fashion. There exists a small literature focusing on the relation between interest and media groups, which tries to address the grassroots aspect of lobbying. Rather than lobbying decision makers, lobby groups aim to influence public opinion and influence decision-making in a direct fashion. The media group acts as an intermediary between interest groups and the population. Interest groups try to influence the population via the media group, which has to decide whether it wants to use biased reports provided by interest groups or higher quality sources for its reporting. Baron (2005) shows that interest groups find it optimal to bias their reports towards their own agenda and profit-maximising media groups find it optimal to bias their reporting towards the

---

[3]This field experiment, where readers were randomly assigned with a prescription to newspapers from the opposite ends of the political spectrum, is cited by some papers as supporting evidence for assuming Bayesian readers. While reading news, in general, increased the probability of voting, which newspaper people were assigned had almost no effect on voting behaviour (everyone's likelihood of voting for the Democratic candidate for governor increased). This is a single field experiment and suffers from some weaknesses; it is not clear why these results would imply that readers were necessarily Bayesian. It may be that the democratic candidate was objectively more competent. In general, the issue is that readers specifically choose which news to follow and seek sources that confirm their beliefs (Klayman, 1995). They might then, after choosing a specific media source, not update beliefs in a Bayesian fashion.

agenda of activist lobbies. Sobbrio (2010), in a more general model incorporating lobbies, media, voters and politicians, shows that interest groups can cause distortions in the political outcome even if they cannot reach the population or decision-makers directly. Li and Mylovanov (2007) attempt to analyse what kind of source a media group uses in an infinite horizon model. The media group can use investigative journalism or reports by interest groups.

The strict distinction between media and interest groups could, in general, be seen as slightly old-fashioned, especially if interest groups use grassroots lobbying. Both types of groups act using the same distribution channels. A model should be able to incorporate them both without assuming some sort of hierarchy by assuming that interest groups somehow need media groups to reach the population. In fact, all groups use similar distribution channels and often cooperate with each other (consider for example the Tea Party and Fox News or the Occupy movement and RT). I will try to take a first step into breaking up the artificial border between the lobbying and media bias literatures by introducing more general objective functions, which go beyond profit-maximisation.

Dynamics

The information media and lobby groups provide can change over time. This could be for several reasons. One obvious reason is change of ownership. A second common reason is that the group has received new information and that its own estimate of the true state of the world has changed. However, I am going to show that media and lobby reports might change over time even if nothing has changed in terms of ownership or the true state of the world. Policies by many interest and media groups have changed, although ownership and the underlying case for their arguments have remained unchanged. In the case of media groups, the general perception seems to be an increasing divide between rival media; for example Fox News shifting further to the right and MSNBC further to the left, abandoning the middle ground.

One can also observe some interest groups becoming more extreme over time, while others seem to have become less extreme. One example of an interest group that has developed a more extreme view over time is the National Rifle Association (NRA), which at least partially supported gun regulation acts in 1934, 1938 and 1968, while it has opposed all recent attempts to regulate gun ownership and has done so with increasingly extreme rhetoric. On the other hand, a group like Greenpeace is nowadays considered almost as a sell-out among many initial members for softening some of its initial policies.

One common explanation for interest groups seemingly becoming more radical over time is that they can 'create their own monsters'. If followers with extreme views do provide an interest group

with stronger support, there exists an incentive for radicalisation, possibly even beyond the own agenda. This theory would suggest that the leaders of the NRA might themselves have a more moderate view, but are catering to their most extreme followers to increase their payoffs. This, in turn, creates a more extreme followership in the following periods. One could create a model explaining this by arguing that those with extreme views do provide stronger support when their views are met by the interest group. This paper, though, provides an alternative explanation. It could be that the interest group itself has extreme views, but rather than catering to its audience it creates its own audience over time. In this setting, the interest group leads the followers towards its own agenda over time. This, as I will show, can explain biases worsening or getting milder over time. Whether biases increase or decrease over time will depend on the extremity of the group's own agenda compared to readers' priors.

While, as mentioned, some of the models in the literature feature Bayesian updating of readers and hint at intertemporal implications (especially Gentzkow and Shapiro, 2006), none of them specifically analyses the development of media biases over time. This paper introduces a basic framework for such an analysis. Li and Mylovanov (2007), as mentioned above, present a dynamic model, but their focus is entirely different as they analyse what kind of sources profit maximising media groups use in optimum.

Effects of competition

In terms of modelling the effects of competition and the general structure of the paper, this paper is closest to Mullainathan and Shleifer (2005). They focus on the effect of confirmation bias and competition on equilibrium bias in a demand side model. Competition has no effect on media biases when followers are homogeneous. However, it has the effect of driving prices down. In the case of heterogeneous readers, each media group becomes more extreme under competition. Individual media biases become worse. A reader following all media groups would be better informed though. Similar results are obtained by Chan and Suen (2008) who show that media groups face an increased incentive to cater to their respective audiences under duopoly. They essentially add electoral competition to a discrete version of Mullainathan and Shleifer (2005).

The results about the effects of competition in the demand side model have been criticised, for example by Gentzkow and Shapiro (2006), because the model lacks reputation effects and the ability of readers to punish the media group. However, while this criticism does not seem invalid, the authors do not provide a comparable model incorporating this.[4] They use a discrete model with

---

[4]Furthermore, there is a cost of slanting built into the model. This cost of slanting, which is modelled to be increasing in slant, should be able to incorporate reputation effects.

low and high quality media and updating by Bayesian readers to show that competition can create different equilibria. Overall, it suggests that competition can reduce media bias, because a media group is more likely to be punished for lying since readers find it easier to fact check given that other media groups exist. Hence, their model resembles the classical economists' view of media. It is a source of information and the more accurate the information, the more valuable it is to readers. However, this relies on all readers actually making use of the ability to fact check reports by comparing them to other reports. Furthermore, the results are driven by the discreteness of the model, meaning that competition cannot increase the wedge between extremes. The model presented in this paper allows for this, as I allow for a continuous, unbounded type space. I investigate the effects of competition on equilibrium bias when demand and supply side motives for bias coexist.

Overall contributions

Most models are limited to discrete states and some of them only allow for a very specific structure of bias, i.e. they allow only for reporting the truth or withholding information and not for simply lying. I try to address this gap with this paper and combine the two different approaches of modelling media biases (demand and supply side). I do not constrain myself to rational readers and allow for a broader set of strategies and forms of confirmation bias. Finally, this paper aims to provide a first step towards establishing a framework that makes it possible to analyse the behaviour of equilibrium bias and ultimately the evolution of public opinion over time.

In doing so, the paper aims to provide answers to the following questions, which remain mainly unanswered by the existing literature:

- How does the coexistence of demand and supply side motives for bias affect equilibrium bias?
- How do existing results on the effects of competition on optimal bias by media or interest groups depend on the relative importance of demand and supply side motives?
- How does equilibrium bias and, therefore, public opinion develop over time?

This paper, by aiming to provide answers to these questions, could help us understand when increased competition can help social cohesion, when it could be harmful and how public opinion could change over time.

2 The Model

2.1 General setup

This paper is not going to model how dispersed public opinion can lead to to a more costly decision-making process. It is also not going to consider how an ill-informed public can lead to bad overall decisions. The model is instead going to analyse how optimal bias provided by media and interest groups depends on priors of followers and on competition. It is also going to provide an insight into how biases might change over time. The assumption that larger biases should, ceteris paribus, be considered bad is merely used for the interpretation of the results. An intuitive explanation for why I use this assumption is provided in the introductory section.

I will take everything that increases slant to be bad in the case of homogeneous readers; the smaller slant or bias, the better. This becomes more complex for the case of heterogeneous readers, because participation constraints play a stronger role. In general, the entire distribution of types has to be regarded. There is no such thing as dispersion of opinions in the homogeneous case. All we want is for the public view to be as close to the true state as possible. However, in the heterogeneous case, we do not just aim for the average view to be as close to the truth as possible, but we actually want little overall variation and deviation from the truth.

Timing

Throughout the paper t refers to time and τ to the true state of the world, which is unknown to readers, but known by the group. It is distributed according to $h(\tau)$. For simplification and without loss of generality: $E(\tau) = 0$ and $E(\tau^2) = v_\tau$. The basic timing of this two-period model is as follows:

$t = 1$

1. The group announces a binding slanting policy $\boldsymbol{s_t(\tau)}$.[5] The binding policy covers slanting policies for both periods, $s_1(\tau)$ and $s_2(\tau)$. If there is more than one group, then all groups announce their slanting policies, $\boldsymbol{s_{t,j}(\tau)}$, simultaneously.

2. Readers decide whether to follow a group and which one to follow. They base this decision on which group provides them with higher expected lifetime utility. Readers make this decision for their lifetime. They are, therefore, not allowed to switch the group or stop following it in $t = 2$. Their expected utility is transferred to the group they follow as support.

---

[5]I will use bold symbols for vectors throughout the paper. $\boldsymbol{s_t(\tau)}$ includes $s_1(\tau)$ and $s_2(\tau)$; $s_t(\tau)$ is not a vector, but only slant at period t.

3. Groups learn about the true state of the world (τ) and publish their reports $\boldsymbol{n_{t,j}(\tau)} = \boldsymbol{s_{t,j}(\tau)} + \tau$. Groups receive their agenda payoff in addition to their support.

4. Reader utilities in $t = 1$ are realised.

$t = 2$

5. Reader i's first period view $b_{1,i}$ is updated to $b_{2,i}$ according to the information he received. This process will be explained in more detail. The groups publish their reports according to their policies they committed to in step 1 and payoffs for the second period are realised.

I assume that readers' participation decisions and the decision which group to follow is based on expected lifetime utility and that that readers remain loyal over both periods. I further assume that groups announce a binding slanting policy for both periods at the beginning of the game. All decisions are, therefore, made in period $t = 1$ (actually in steps 1 and 2 above). The relevant decisions are made by consumers based on binding slanting policies $\boldsymbol{s_{t,j}(\tau)}$. Optimisations are made at the ex ante stage. Steps 3, 4 and 5 are only technical steps and do not involve any decision-making by agents. This setup is chosen for simplification and aims to capture the idea that this is really a long term game. Readers are aware of the general positioning of different news outlets, i.e. they know which one tends to be conservative and which one liberal etc., but they do not know exact reports before deciding to follow a specific media group.

Followers/readers

Followers dislike slant because finding out the truth comes with costly effort. When news reports are slanted, readers could deduce the true value of τ after observing media slant, but this process is a hassle. The parameter χ aims to capture this cost, which is increasing in the size of slant. Throughout the paper, I assume that all readers dislike slant equally.[6,7]

Followers also suffer from confirmation bias. They dislike hearing information that does not match their current views. The private view of follower i in period t is $b_{t,i}$, which is distributed according to $f_t(b_{t,i})$. The parameter $\phi(b_{t,i})$ measures the strength of confirmation bias and is allowed to depend on the reader's type in the most general setting.

Private views at period t and their overall distribution, which I refer to as public opinion, will change

---

[6]I allowed for $\chi(b_{t,i})$ in an earlier version. It does not really have any meaningful effect in this model, because there is no relation between it and readers' priors. Such a relationship exists for the confirmation bias $\phi(b_{t,i})$, which is why I will focus more on confirmation bias in this paper.

[7]All that is necessary for this model is that there is some cost of slanting to the media group; it could also be modelled in any other way. One could introduce a mechanism that stops readers from being able to deduce τ. This could be achieved by making media agenda private information and introducing a verification stage in which readers can find out whether they have been lied to by the media (the probability of which would be increasing in the size of slant). If readers find out that they have been lied to, they would then punish the media group.

over time. I will not model the evolution of followers' views as the result of some maximisation problem. Followers, in this model, do not update their beliefs in a Bayesian fashion. The degree to which followers use the information provided by the media or interest group is not the result of a dynamic utility maximisation problem. Instead, the evolution of private opinions is assumed to be mechanical and happens according to the following rule:

$$b_{2,i} = \begin{cases} \gamma n_{1,j}(\tau) + (1-\gamma)b_{1,i} & \text{if } EU_{i,j} > EU_{i,k}, \ EU_{i,j} \geq 0 \\ b_{1,i} & \text{if } EU_{i,j} > EU_{i,k}, \ EU_{i,j} < 0 \end{cases}$$

where $EU_{i,j}$ is the expected lifetime utility for reader i from following group j. The precise nature of expected lifetime utility is explained below.

Hence, an agent i follows the group that provides him with higher lifetime utility and only participates at all when this utility is weakly positive. Otherwise, the agent opts to not follow any group and his views do not change over time; in other words, γ is zero for agents whose participation constraint is violated. The assumption that participation choices and the choice about which media to group to follow are done in period one, and cannot be changed, is made for simplification. Given the rest of the model setup, it seems like the only possible equilibrium outcome anyways. There is nothing within the model that would incentivise followers to switch the group they are following.

The updating assumption above drives some of the conclusions of the paper, especially any results with regards to how biases change over time, because the evolution of private opinions is assumed rather than derived endogenously. This could, in future research, be achieved by making the parameter γ a choice variable in a utility maximisation problem of followers.

The parameter γ, even if not endogenous, could be assumed to depend on the type or on the distance between own belief and the group's report in a more flexible model. I am going to assume it to be constant throughout this paper for simplicity. It seems to make intuitive sense to speculate that reports would be used less for updating if the report did not match the agent's prior (this is what the literature on confirmation bias suggests).[8] However, there would also be no need for such agents to use the report heavily, because it deviates from their views only slightly. Overall, I do not want to speculate on this too much and want to keep the model as simple as model. This element is also already captured partially by the parameter ϕ, which I will allow to be type-dependent.

Expected lifetime utility $EU_{i,j}$ is just the sum of expected utilities in each period. It will, thus, depend on the slanting policy of the media group j that the agent decides to follow $(s_{t,j}(\tau))$ and the private views of the follower in both periods, which is summarised by $\boldsymbol{b_{t,i}}$.[9]

---

[8]The psychology literature on confirmation bias suggests this to be the case and provides evidence that the existence of confirmation bias can, therefore, lead to increased polarisation over time (Kuhn and Lao, 1996). The model presented here will derive similar results.

[9]Once again, I use bold symbols for vectors. $\boldsymbol{b_{t,i}}$ includes $b_{1,i}$ and $b_{2,i}$.

Reader i's expected utility from following media group j is, thus, given by:

$$EU_{i,j}(\boldsymbol{s_{t,j}(\tau)}, \boldsymbol{b_{t,i}}) = Eu_{i,j}(s_{1,j}(\tau), b_{1,i}) + Eu_{i,j}(s_{2,j}(\tau), b_{2,i})$$

$$= \sum_{t=1}^{2} \left(\frac{\overline{u}}{2} - \chi \int_{\tau} s_{t,j}(\tau)^2 h(\tau) d\tau - \phi(b_{t,i}) \int_{\tau} (s_{t,j}(\tau) + \tau - b_{t,i})^2 h(\tau) d\tau \right).$$

I will not use any discounting for simplicity, because the model already has a number of parameters. I will also focus on the simple two period setting throughout the paper. This is sufficient to outline the basic principles. A dynamic programming approach might be promising for the future.

The parameter $\overline{u}$ is necessary for determining the follower's participation constraint:

$$EU_{i,j}(\boldsymbol{s_{t,j}(\tau)}, \boldsymbol{b_{t,i}}) \geq 0.$$

Only agents with positive expected lifetime utility will opt to follow a group at all, while agents deriving a negative expected lifetime utility will opt out. As already mentioned, I assume that all agents with $EU_{i,j}(\boldsymbol{s_{t,j}(\tau)}, \boldsymbol{b_{t,i}})$ will participate in both periods. I do not allow agents to follow a group in the first period and drop out in the second period. I, furthermore, assume that readers remain loyal to the group they follow.

Media/interest group

The group maximises a mix of readers' expected utilities and own direct benefit derived from spreading its agenda among readers. When an agent follows a specific group, his expected utility is transferred to it. The group places a weight ω on the support (transferred lifetime utility) and $1 - \omega$ on spreading its agenda.

Another difference to other models, for example Mullainathan and Shleifer (2005), is that the reports are not sold at a price. Agents, in the model presented here, follow a specific group and groups do not necessarily want to maximise the amount of followers, but they also want to create strong followers. They do this not because they can sell reports at a higher price, but because utility transfers directly.[10]

The objective function can be written either in terms of $\boldsymbol{s_{t,j}(\tau)}$ or $\boldsymbol{n_{t,j}(\tau)}$, since they are solely dependent on each other. I will model the choice to be slant, but results are independent of this.

---

[10]I think this captures the objective of modern media and lobby groups better than a classic profit maximisation problem. In a classic profit maximisation model, all readers pay the same and, hence, have the same intrinsic value to the media group. This model has a sort of probabilistic voting model concept. The more utility the reader derives from following a specific media/lobby group the more benefit he provides to it in return; it is as if the media group was able to use price discrimination. In a model with prices, it is only the utility of marginal readers that matters, because it determines the price for everyone. In the model presented here, like in a model of price discrimination, the total sum of utilities is what matters. As an example, this could be because a lot of revenue comes via advertisement and advertisers would be willing to pay much more for strongly engaged followers. When thinking about lobby groups, a small group of heavily motivated followers usually is much stronger in creating political influence than large groups consisting of only slightly engaged members.

Write group j's objective function Π_j in its most general form as:

$$\Pi_j(\boldsymbol{s_{t,j}(\tau)}) = \sum_t \pi_{t,j}(s_{t,j}(\tau))$$

$$= \sum_t \left(\int_{b_{1,i}^{j,min}(\cdot)}^{b_{1,i}^{j,max}(\cdot)} \left(\omega E u_{i,j}(s_{t,j}(\tau), b_{t,i}) + (1-\omega)A_t(s_{t,j}(\tau)) \right) f_t(b_{t,i})db_{t,i} \right)$$

where the thresholds $b_{1,i}^{j,min}(\cdot)$ and $b_{1,i}^{j,max}(\cdot)$ depend on the group's slanting policy $\boldsymbol{s_{t,j}(\tau)}$ and

$$A_t(s_{t,j}(\tau)) = \frac{\overline{a}}{2} - (s_{t,j}(\tau) + \tau - a_j)^2$$

is the payoff the group derives in period t from spreading its agenda in that period. It is modelled as a loss function and the purely technical parameter $\overline{a}$ is needed to ensure payoffs are positive; otherwise, increasing the number of followers would enter the payoff function negatively, which would not match the story and intuition. The point about the agenda payoff is that the media likes to spread it as widely as possible. The groups dislike to publish reports $s_{t,j}(\tau) + \tau$ that are different to their own agenda a_j. The media group does not care who the readers are intrinsically, i.e. the agenda payoff is just proportional to the total number of readers. Hence, there is no participation constraint for the media/lobby group by assumption. Groups always benefit from increasing their followership given a fixed news position. This means that the media/lobby group will serve a follower and will never choose not to serve a follower who wants to be served. Whenever the reader's participation constraint holds, so will the group's constraint. The thresholds $b_{1,i}^{j,min}(\boldsymbol{s_{t,j}(\tau)})$ and $b_{1,i}^{j,max}(\boldsymbol{s_{t,j}(\tau)})$ are identical in period $t = 2$, because followers are assumed to make all participation decisions in period $t = 1$ and to be loyal. All types between the two thresholds are followers of media group j.

In summary, the group chooses a slant taking into account which readers will follow it given the slant. By doing this, it maximises a weighted sum of the overall sum of followers' utilities, which transfers to it as support, and its own agenda. The distinction between a media group and a lobby or interest group can be summed up in the value of ω. A traditional media group has a stronger profit driven agenda and wants a large readership. This would mean it has a high ω. An ω close to one would resemble a pure profit maximiser with the caveats explained above. A lobby group, on the other hand, has a strong agenda and tries to build a supportive group with fellow beliefs, or tries to convince people of its own agenda. It is, thus, best described by a low ω.

2.2 Homogeneous agents

Consider the simplest case of homogeneous agents to better understand the workings of the model. Structuring the paper in the same way as Mullainathan and Shleifer (2005) also allows us to compare directly the results of this intertemporal model with groups having an own agenda to their static model without agenda. Furthermore, views about specific issues, such as attitudes towards other societies/countries or social insurance schemes, are actually often not very heterogeneous in many societies. Although opinions will never be completely homogeneous, this setting works as a useful benchmark and will later on help to explain how heterogeneity affects the results. In this simpler case $b_{1,i} = b_1$, $b_{2,i} = b_2$ and $\phi(b_{t,i}) = \phi$.

2.2.1 Monopoly

Slant and types are not indexed by j or i, because there is only a single group and agents are homogeneous. There are no thresholds b_1^{min} and b_1^{max} as there is only one type of agent, who either follows the group or not. The simplified utility function of the agent is, thus, given by:

$$EU_{hom} = \bar{u} - \chi \int_\tau s_1^2 h(\tau)d\tau - \phi \int_\tau (s_1 + \tau - b_1)^2 h(\tau)d\tau - \chi \int_\tau s_2^2 h(\tau)d\tau - \phi \int_\tau (s_2 + \tau - b_2)^2 h(\tau)d\tau.$$

The simplified objective function, which the group aims to maximise from the ex ante perspective, is given by:

$$\Pi_{hom} = \omega EU_{hom} + (1 - \omega) \int_\tau \left(\bar{a} - (s_1 + \tau - a)^2 - (s_2 + \tau - a)^2 \right) h(\tau)d\tau.$$

Proposition 1 (Monopolistic slant with homogeneous readers)**.**

As long as $EU(\boldsymbol{s_t^*(\tau)}, b_1) \geq 0$, the monopolist's optimal slanting policy is given by:

$$s_1^*(\tau) = \alpha_1^*(a - \tau) + \beta_1^*(b_1 - \tau)$$

$$s_2^*(\tau) = \alpha_2^*(a - \tau) + \beta_2^*(b_1 - \tau)$$

where

$$\beta_1^* = \frac{\phi\omega(1 + (\phi + \chi - 1)\omega + \gamma^2(1 + (\chi - 1)\omega) + \gamma(-1 + \omega - \chi\omega))}{1 + ((2 + \gamma^2)\phi + 2(\chi - 1))\omega + (\phi^2 + (2 + \gamma^2)\phi(\chi - 1) + (\chi - 1)^2)\omega^2}$$

$$\alpha_1^* = \frac{(1 - \omega)(1 + (\phi + \gamma\phi + \chi - 1)\omega)}{1 + ((2 + \gamma^2)\phi + 2(\chi - 1))\omega + (\phi^2 + (2 + \gamma^2)\phi(\chi - 1) + (\chi - 1)^2)\omega^2}$$

$$\beta_2^* = \frac{\phi\omega(1 + (\phi + \chi - 1)\omega + \gamma(-1 + \omega - \chi\omega))}{1 + ((2 + \gamma^2)\phi + 2(\chi - 1))\omega + (\phi^2 + (2 + \gamma^2)\phi(\chi - 1) + (\chi - 1)^2)\omega^2}$$

$$\alpha_2^* = \frac{(1 - \omega)(1 + (\phi + \gamma\phi + \gamma^2\phi + \chi - 1)\omega)}{1 + ((2 + \gamma^2)\phi + 2(\chi - 1))\omega + (\phi^2 + (2 + \gamma^2)\phi(\chi - 1) + (\chi - 1)^2)\omega^2}.$$

Proof. Proofs for propositions are provided in appendix B, supporting lemmata in appendix A. The exception to this is proposition 7 which is derived in the main body of the paper; it follows directly from the analysis preceding it. □

This proposition defining optimal slant with homogeneous readers is the equivalent to propositions 1 and 2 in Mullainathan and Shleifer (2005). It extends the proposition in that paper by introducing simple dynamics and allowing for a supply side motive of slant. The group slants around two points: initial prior (b_1) and own agenda (a) with weights $\boldsymbol{\beta_t^*}$ and $\boldsymbol{\alpha_t^*}$. It is shown easily that $1 \geq \beta_1^*, \alpha_1^*, \beta_2^*, \alpha_2^* \geq 0$. The size of slant is, therefore, determined by the distance of these focal points from the true state and by the magnitudes of the weights, which depend on the parameters of the model (ω, ϕ, χ and γ). There are two sources of bias in reporting: confirmation bias and own agenda.

There will be slant even when there is no confirmation bias unless the group does not care about own agenda at all ($\omega = 1$) or the agenda is aligned with the truth ($a = \tau$). Proposition 1 makes clear how the group trades off demand side motives (such as in Mullainathan and Shleifer, 2005) and supply side motives (such as in Anderson and McLaren, 2010). It shows how the size and direction of overall bias depends on both. When both motives have the same direction, bias will lie between the two motives. It is, however, possible that the two motives can eliminate each other, creating small overall bias.

Having established what determines slant, we can focus on how it is affected by the underlying parameters of the model. Taking simple derivatives, we can see that the weight that the group places on utility transfer versus own agenda affects the weights of optimal slant as follows:

$$\frac{\partial \beta_1^*}{\partial \omega} \geq 0, \frac{\partial \alpha_1^*}{\partial \omega} \leq 0; \quad \frac{\partial \beta_2^*}{\partial \omega} \geq 0, \frac{\partial \alpha_2^*}{\partial \omega} \leq 0$$

$$\left| \frac{\partial \beta_1^*}{\partial \omega} \right| \leq \left| \frac{\partial \alpha_1^*}{\partial \omega} \right|; \quad \left| \frac{\partial \beta_2^*}{\partial \omega} \right| \leq \left| \frac{\partial \beta_2^*}{\partial \omega} \right|.$$

The size of confirmation bias affects the weights in the following way:

$$\frac{\partial \beta_1^*}{\partial \phi} \geq 0, \frac{\partial \alpha_1^*}{\partial \phi} \leq 0; \quad \frac{\partial \beta_2^*}{\partial \phi} \geq 0, \frac{\partial \alpha_2^*}{\partial \phi} \leq 0$$

$$\left| \frac{\partial \beta_1^*}{\partial \phi} \right| \geq \left| \frac{\partial \alpha_1^*}{\partial \phi} \right|; \quad \left| \frac{\partial \beta_2^*}{\partial \phi} \right| \geq \left| \frac{\partial \alpha_2^*}{\partial \phi} \right|.$$

Dislike of slant has the following effects:

$$\frac{\partial \beta_1^*}{\partial \chi} \leq 0, \frac{\partial \alpha_1^*}{\partial \chi} \leq 0; \quad \frac{\partial \beta_2^*}{\partial \chi} \leq 0, \frac{\partial \alpha_2^*}{\partial \chi} \leq 0.$$

Finally, the extent to which the readers use news reports to update their views affects the weights

in the following way:

$$\frac{\partial \beta_1^*}{\partial \gamma} \geq 0 \text{ if } \gamma \geq \overline{\gamma}, \frac{\partial \beta_1^*}{\partial \gamma} \leq 0 \text{ if } \gamma \geq \overline{\gamma}; \quad \frac{\partial \beta_2^*}{\partial \gamma} \leq 0, \frac{\partial \alpha_2^*}{\partial \gamma} \geq 0$$

$$\left| \frac{\partial \beta_1^*}{\partial \gamma} \right| \geq \left| \frac{\partial \alpha_1^*}{\partial \gamma} \right|; \left| \frac{\partial \beta_2^*}{\partial \gamma} \right| \geq \left| \frac{\partial \alpha_2^*}{\partial \gamma} \right|.$$

Combining and summarising all of these partial effects leads to the following corollary, which summarises how slant is affected by the underlying parameters of the model.

Corollary 1 (How does monopolistic slant depend on underlying parameters?).

- *An increase in ω causes an increase in slant only if the group's agenda is sufficiently more extreme than the follower's prior.*

- *An increase in ϕ causes a decrease in slant only if the group's agenda is sufficiently less extreme than the follower's prior.*

- *An increase in χ reduces both weights and, therefore, implicitly increases the weight on truth.*

- *An increase in γ causes adverse effects on the weights and potentially different effects in terms of direction in the first and second period depending on the initial level on γ. Overall, it causes a decrease in slant in either period only if the group's agenda is sufficiently less extreme than the follower's prior.[11]*

Introducing a dynamic approach and own media agenda does, thus, not change the influence of χ on bias. It does affect the influence of ϕ, since the influence of confirmation bias now depends on the relative positions of priors and group agenda. It is not clearly determined as in Mullainathan and Shleifer (2005). This is also true for the influence of the newly introduced parameters ω and γ.

Finally, I want to put special focus on how slant and public opinion develop over time. This aspect is missing from the existing literature. Is there a pressure on slant to worsen over time and for public opinion to diverge from the truth, or should public opinion eventually converge to the true state when it can be manipulated with biased reports?

Proposition 2 (Slant over time).

Slant is decreasing over time in a symmetric setting $((|a - \tau| = |b_1 - \tau|))$. Slant is increasing over time when the agenda is sufficiently more extreme than the reader's prior. There is a wider set for the agenda parameter under which slant is decreasing over time the higher χ, ω and ϕ.

---

[11]By sufficiently more or less extreme I mean that it would not happen in a symmetric setting, by which I mean $|a - \tau| = |b_1 - \tau|$. For example, a sufficiently more extreme agenda means that $|a - \tau| > |b_1 - \tau|$ by some amount that depends on the parameters of the model.

We can summarise and visualise the two conditions and proposition 2 by building on the proof provided in the appendix. We can show how the conditions and, therefore, slant over time depend on the underlying parameters in the following graph:

Figure 1: Slant over time
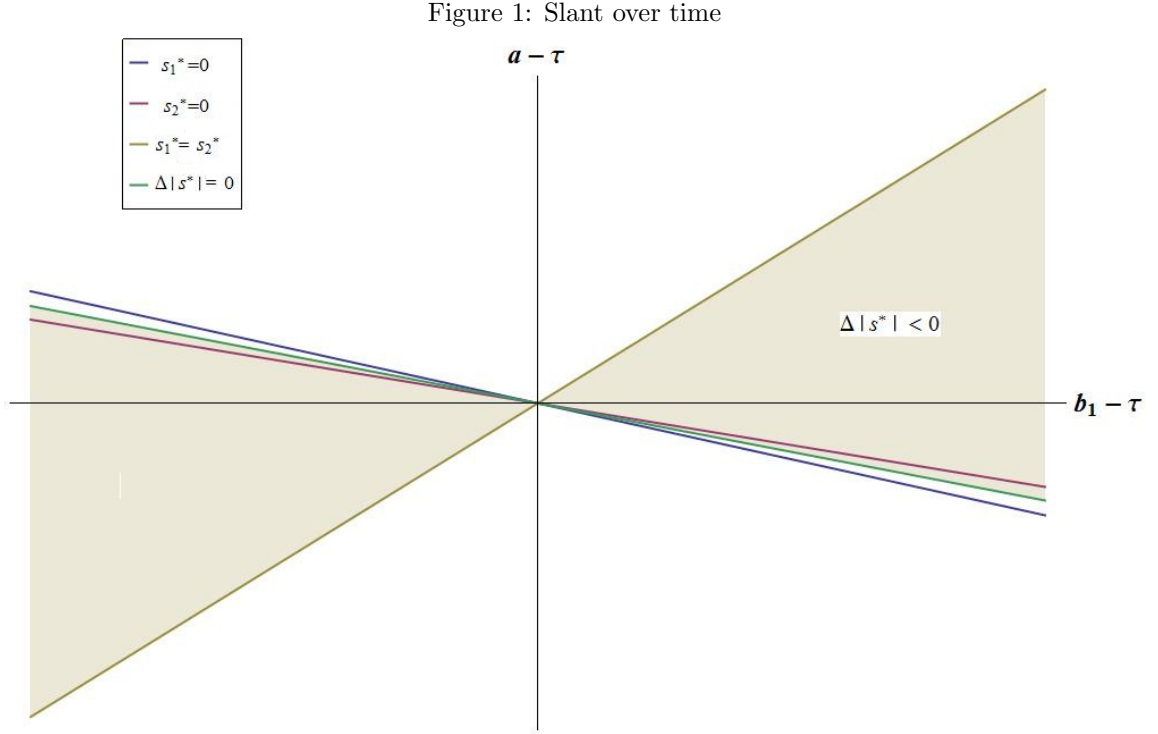
All points above the blue line are points in which first period slant is positive. Points below the line are combinations of priors and agenda for which first period slant is negative. The same holds for second period slant and the red line. All points above the golden line are points at which second period slant is larger than first period slant, while all points below are combinations of priors and agendas for which second period slant is smaller. This implies that the area below the golden and above the red line are combinations of priors and agenda for which slant, in absolute terms, is clearly decreasing over time as long as $b_1 - \tau \geq 0$. If $b_1 - \tau < 0$, then this is true for the area below the red, but above the golden line. Areas above the blue and golden lines are combinations for which slant worsens over time for $a - \tau \geq 0$, while the same holds for the area below the blue and golden lines if $a - \tau < 0$. What happens in the area between the red and blue lines is unclear, because the sign of slant switches between the two periods. In general, there will be an area, indicated by the green line in the graph, that is determined by $|s_2^*(\tau)| = |s_1^*(\tau)|$. It will necessarily lie between the blue and red line with its exact position depending on the parameters of the model. Combining all of this, we can say that there exists a set of agendas and priors for which slant gets better over

time (here shown by the coloured area) and a set of agendas and priors for which slant worsens (the non-coloured area in this graph).

2.2.2 Duopoly

I am going to focus on a simple duopoly to analyse competition. This gives us insight into how competition affects the market in the simplest fashion. The expected lifetime utility of the agent must be equal for following either of the groups in equilibrium. Furthermore, it has to be maximised. Otherwise, a group could offer a slanting policy that increases the agent's utility and attract the entire market.

Proposition 3 (Duopolistic slant with homogeneous readers).

As long as $EU(\boldsymbol{s_t^P(\tau)}, b_1) \geq 0$, the optimal slanting policy is identical for both duopolists and given by:

$$s_1^D(\tau) = \beta_1^D(b_1 - \tau)$$

$$s_2^D(\tau) = \beta_2^D(b_1 - \tau)$$

where

$$\beta_1^D = \frac{\phi(\phi + (1 - \gamma(1 - \gamma))\chi)}{\phi^2 + (2 + \gamma^2)\phi\chi + \chi^2}$$

$$\beta_2^D = \frac{\phi(\phi + (1 - \gamma)\chi)}{\phi^2 + (2 + \gamma^2)\phi\chi + \chi^2}$$

and implicitly

$$\alpha_1^D = 0$$

$$\alpha_2^D = 0.^{12}$$

Competitive slant is unambiguously decreasing over time: $|s_1^D(\tau)| > |s_2^D(\tau)|$.

Competition eliminates the emphasis on agenda by groups. The outcome, in terms of slant, is identical to a single firm with $\omega = 1$. A duopoly is sufficient to induce Bertrand-like competition and eliminate slant driven by group agenda completely, even though there is no such thing as price competition in the model. In a model with pricing and no supply side motive for slanting, Mullainathan and Shleifer (2005) show that the slants under monopoly and duopoly are identical and that only prices are driven down by competitive pressure. Hence, competition works by driving prices down. I have shown here that the supply side motive is eliminated in competition and that

---

[12]Note, that the outcome would be the same for any number of groups. In a sense, although there is no pricing in this model, groups face Bertrand competition.

the weight on the demand side motive is decreased as well ($\beta_t^D < \beta_t^*$) in a model of transferable utility rather than pricing.

The effects of underlying parameters on the weights in the slanting policy are given by the following:

$$\frac{\partial \beta_1^D}{\partial \omega} = 0; \quad \frac{\partial \beta_2^D}{\partial \omega} = 0$$

$$\frac{\partial \beta_1^D}{\partial \phi} \geq 0; \quad \frac{\partial \beta_2^D}{\partial \phi} \geq 0$$

$$\frac{\partial \beta_1^D}{\partial \chi} \leq 0; \quad \frac{\partial \beta_2^D}{\partial \chi} \leq 0$$

$$\frac{\partial \beta_1^D}{\partial \gamma} \geq 0 \text{ if } \gamma \geq \overline{\gamma}^D; \quad \frac{\partial \beta_2^*}{\partial \gamma} \leq 0.$$

These are summarised in corollary 2.

Corollary 2 (Slant in competition with homogeneous agents)**.**

First period slant is unambiguously increasing in ϕ and increasing in γ if γ is sufficiently large, while it is decreasing in χ and independent of ω. Second period slant is increasing in ϕ, decreasing in χ and γ, while it is independent of ω.

Having established equilibrium slant, how it develops over time and how it depends on underlying parameters in a monopolistic as well as a competitive setting, we can now focus on comparing the two.

Proposition 4 (Comparison of slant under monopoly and competition)**.**
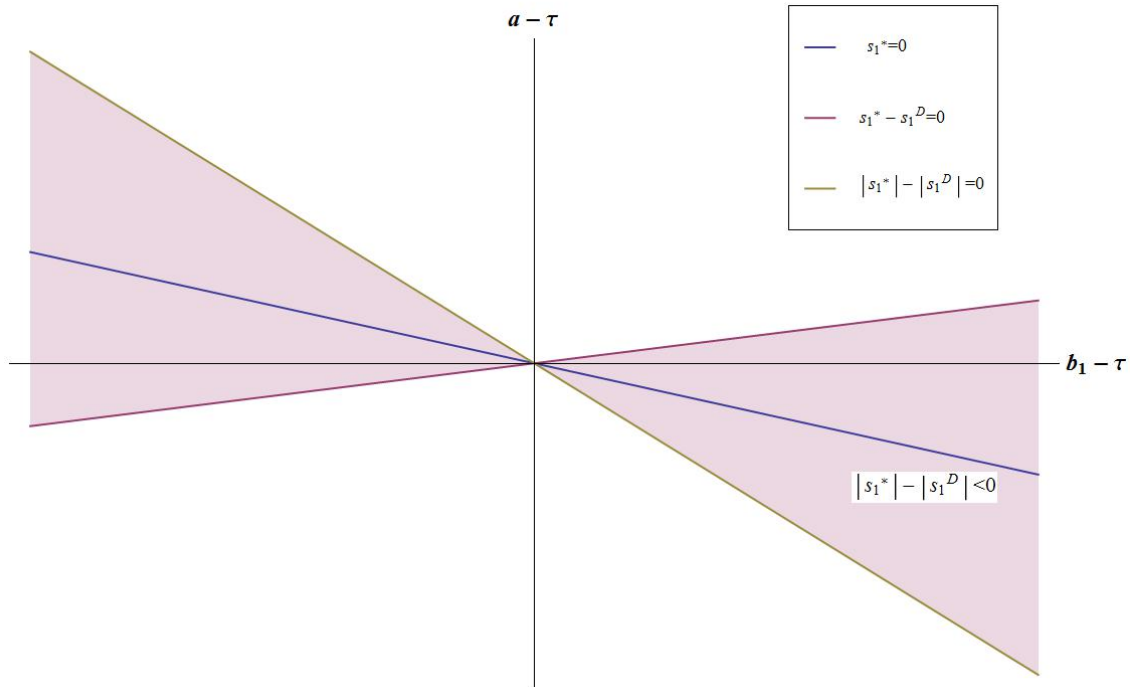
- *There exists a set of agendas for which slant (in absolute terms) is smaller under monopoly than it is under competition.*
- *The set of agendas for which absolute slant is smaller (in absolute term) under monopoly decreases over time.*

Essentially, there are two sources of slant: confirmation bias/agents' priors and group agenda. The agenda motive is eliminated under competition and the combined weight assigned in the slanting policy to agents' priors and the truth is, therefore, larger. If agents' priors are extreme and far off the true state, this can imply large slants. If the agenda is more modest or even goes into the opposite direction of priors, it can weaken or partially offset the slant driven by confirmation bias in the monopolistic case. Equilibrium slants can then be larger under competition in these cases.

Over time, though, the weight assigned to agents' priors decreases and the weight on the truth increases. The homogeneous view of agents converges towards the true state under competition, because agents dislike slant and their opinions evolve over time. It follows that competition becomes more favourable over time unless the agenda is aligned with the true state, because the convergence of public opinion under monopoly is towards a point determined by a combination of agenda and the truth and not towards the truth alone.

We can visualise the proposition in the following graph. It summarises how slant under competition compares to monopolistic slant.

Figure 2: Monopolistic versus competitive slant

All points above the blue line are combinations of agendas and priors for which first period monopolistic slant is positive. Monopolistic slant is negative at points below that line. First period and second period slants under competition are positive at all points to the right of the y-axis, i.e. if $b_1 > \tau$. All points above the red line are combinations for which first period slant is larger under monopoly than under competition. The area between the red and blue lines, therefore, depicts the set of agendas for which slant, given a prior, is definitely smaller under monopoly in absolute terms. The opposite is true for the area to the right of the y-axis and the red line for $a - \tau \geq 0$ and to the left of the y-axis and the red line if $a - \tau < 0$. The other area remains unclear. Like

in the analysis of slant over time, there will be a cut-off visualised by the golden line that depicts $|s_1^*(\tau)| = |s_1^D(\tau)|$, so that the entire coloured area are combinations of slants and agendas for which absolute slant is lower under monopolistic provision. All lines become flatter in the second period; hence, decreasing the set favouring monopoly.

2.3 Heterogeneous readers

2.3.1 Monopoly

While I allow for a distribution of reader views, I still maintain $\phi(b_{t,i}) = \phi$. With participation decisions being made in the first period, the group's objective function is given by:

$$\Pi = \int_{b_1^{min}(\boldsymbol{s_t}(\boldsymbol{\tau}))}^{b_1^{max}(\boldsymbol{s_t}(\boldsymbol{\tau}))} \left(\omega EU(\boldsymbol{s_t}(\boldsymbol{\tau}), b_{1,i}) + (1-\omega)\Big(A_1(s_1(\tau)) + A_2(s_2(\tau))\Big) \right) f(b_{1,i}) db_{1,i}.$$

Once again, the objective function and underlying participation decision of followers can be written as solely depending on $b_{1,i}$, because binding participation decisions are made at $t = 1$. Remembering that $\pi_t = \omega Eu(s_t(\tau)) + (1-\omega)A_t$, i.e. that the group's payoff is the sum of transferred expected utility and agenda payoff in each period, and noting that $b_1^{max}(\boldsymbol{s_t}(\boldsymbol{\tau}))$ and $b_1^{min}(\boldsymbol{s_t}(\boldsymbol{\tau}))$ are defined by a binding participation constraint of followers and, hence, $EU(b_1^{min}(\cdot)) = EU(b_1^{max}(\cdot)) = 0$, we can say that the group's payoff for marginal followers is just given by agenda payoffs, which is independent of follower prior: $\Pi(b_1^{min}(\cdot)) = \Pi(b_1^{max}(\cdot)) = A_1(s_t) + A_2(s_t)$. This leads to the following first order conditions, which when solved yield optimal slant $\boldsymbol{s_t^{het}}(\boldsymbol{\tau})$:

$$\frac{\partial \Pi}{\partial s_1} = \int_{b_1^{min}(\boldsymbol{s_t}(\boldsymbol{\tau}))}^{b_1^{max}(\boldsymbol{s_t}(\boldsymbol{\tau}))} \left(\omega \frac{\partial EU(\boldsymbol{s_t}(\boldsymbol{\tau}), b_{1,i})}{\partial s_1} + (1-\omega)\frac{\partial A_1(s_1(\tau))}{\partial s_1} \right) f(b_{1,i}) db_{1,i}$$

$$+ \Big(A_1(s_1(\tau)) + A_2(s_2(\tau))\Big) \left(f(b_1^{max}(\boldsymbol{s_t}(\boldsymbol{\tau})))\frac{\partial b_1^{max}(\boldsymbol{s_t}(\boldsymbol{\tau}))}{\partial s_1} - f(b_1^{min}(\boldsymbol{s_t}(\boldsymbol{\tau})))\frac{\partial b_1^{min}(\boldsymbol{s_t}(\boldsymbol{\tau}))}{\partial s_1} \right)$$

$$\frac{\partial \Pi}{\partial s_2} = \int_{b_1^{min}(\boldsymbol{s_t}(\boldsymbol{\tau}))}^{b_1^{max}(\boldsymbol{s_t}(\boldsymbol{\tau}))} \left(\omega \frac{\partial EU(\boldsymbol{s_t}(\boldsymbol{\tau}), b_{1,i})}{\partial s_2} + (1-\omega)\frac{\partial A_2(s_2(\tau))}{\partial s_2} \right) f(b_{1,i}) db_{1,i}$$

$$+ \Big(A_1(s_1(\tau)) + A_2(s_2(\tau))\Big) \left(f(b_1^{max}(\boldsymbol{s_t}(\boldsymbol{\tau})))\frac{\partial b_1^{max}(\boldsymbol{s_t}(\boldsymbol{\tau}))}{\partial s_2} - f(b_1^{min}(\boldsymbol{s_t}(\boldsymbol{\tau})))\frac{\partial b_1^{min}(\boldsymbol{s_t}(\boldsymbol{\tau}))}{\partial s_2} \right).$$

We can see the main difference to the homogeneous case in the second lines of the first order conditions. The group has to take into account participation constraints. The thresholds $b_1^{min}(\boldsymbol{s_t}(\boldsymbol{\tau}))$ and $b_1^{max}(\boldsymbol{s_t}(\boldsymbol{\tau}))$ depend on the slanting policy. The group has to take into account the marginal

benefits it gets from the marginal readers at both ends of the spectrum.

Proposition 5 (Monopolistic slant with heterogeneous readers)**.**

The monopolist's slanting policy is given by:

$$s_t(\tau)^{het} = \alpha_t^*(a - \tau) + \beta_t^*(b_1^{het} - \tau).$$

If $b_{1,i}$ is symmetrically distributed around $\tilde{b}_1$, then:

$$b_1^{het} = \bar{b}_1 = \tilde{b}_1 = \hat{b}_1.$$

If $b_{1,i}$ is not symmetrically distributed around $\tilde{b}_1$, then $b_1^{het} = \hat{b}_1$ lies between $\bar{b}_1$ and $\tilde{b}_1$.

Since the weights are the same as in the homogeneous case, any results about how slant develops over time carry through to the heterogeneous case.

As a reminder and to understand the proposition and its intuition, remember that:

- $\bar{b}_1 = E(b_{1,i})$ is the expected, and with a unit mass, average prior of all followers in the entire population;

- $\hat{b}_1$ is the expected prior of those followers who actually follow the group; and

- $\tilde{b}_1$ is defined in lemma 1. It is the reader who is most valuable in terms of support when having his prior matched by the group.

Essentially, if it happens to be the case that the distribution is such that $b_{1,i}$ is symmetrically distributed around $\tilde{b}_1$, the group can maximise its payoff as if the marginal readers were independent of the slanting policy, because marginal gains and losses from marginal followers cancel out. The group can, therefore, slant towards the average follower's prior (as in the homogeneous case). Otherwise, the group will skew slant towards the marginal reader that provides the group with larger benefit. This might be the case because of a skewed distribution or because of the average prior not matching $\tilde{b}_1$.

This result resembles and extends the result of Mullainathan and Shleifer (2005), which is the special case of $a = 0$ and/or $\omega = 1$ as well as $\gamma = 0$. This shows that results are not affected qualitatively whether one explicitly uses pricing or not . Quantitative differences stem from the fact that, unlike in Mullainathan and Shleifer (2005), $\tilde{b}_1 \neq 0$ in the dynamic setting when groups have an own agenda.

2.3.2 Monopoly and type-dependent confirmation bias

In general, I could allow for ϕ to depend on current types ($\phi(b_{t,i})$), so that confirmation bias changes in tune with agents' views. I am instead, for simplification purposes, going to assume that these parameters are determined in the first period. Confirmation biases differ across agents, but not over time for each agent. This simplifies the model greatly. Furthermore, this assumption will only affect results quantitatively and not qualitatively, because what matters is how parameters compare across agents. There will be no change in order with γ being constant, although followers' views will change over time ($b_{1,i} \geq b_{1,j}$ implies $b_{2,i} \geq b_{2,j}$ and vice versa).

Allowing for type-dependency of such kind ($\phi(b_{1,i})$), we can derive the group's first order conditions which, when solved, will determine optimal slant with type-dependent confirmation bias (s_t^{TD}). The first order conditions look identical to the ones presented in the case of constant confirmation bias. The only difference is that the average confirmation bias of followers is now a function of slant. The slant provided determines the marginal readers which in turn determine the average confirmation bias the group faces.

Proposition 6 (Monopolistic slant with type-dependent confirmation bias)**.**

A group providing positive slant (negative slant) faces an incentive to increase (decrease) slant compared to the scenario in which confirmation bias is not type-dependent at the very least if $\phi(b_{1,i})$ is symmetrically distributed around $b_{1,i} = 0$, increasing in $|b_{1,i}|$ and own agenda is not too extreme compared to b_1^{het}. This means that a group increases bias in absolute terms in this setting. This is true unless $s_t(\tau)^{het} = 0$.

When confirmation bias is increasing in the extremity of followers' priors, the group is getting 'punished' in a disproportionate way by its more extreme followers for not matching their priors. The group's optimal reaction tends to be to appease these followers. It trades off lost support of moderate followers with gained support of more extreme followers. The support of more extreme followers becomes relatively more valuable. This effect can be especially large when the group puts a very strong emphasis on support rather than own agenda.

2.3.3 Duopoly

I am going to focus on the case of a duopoly, because it gives us significant insight into how competition affects the flow of information while keeping the model as simple as possible. Market

shares remain constant over time due to our assumption of loyalty.[13]

In period $t = 1$ reader i has to decide which group to follow. Groups have agendas of a_1 and a_2, where I assume $a_1 < a_2$. As both groups face the same overall distribution of agents, this necessarily implies that:

$$s_{t,1}(\tau)^{het,D} < s_{t,2}(\tau)^{het,D} \quad \text{for } t = 1, 2.[14]$$

This means that a group with a left-wing agenda is necessarily going to have a slanting policy and, hence, a bias that is more left-wing than the slanting policy and bias of a right-wing group. I am interested in whether these slanting policies are more extreme under competition than they would be without it and in how they develop over time. Two scenarios are conceivable under duopoly.

The first scenario is that optimal slanting policies of groups are so extreme that no reader is actually indifferent between following either group. The market size for each group is then just given by the followers' participation constraints; all four of which are binding. This is depicted in figure 3.

Figure 3: scenario 1
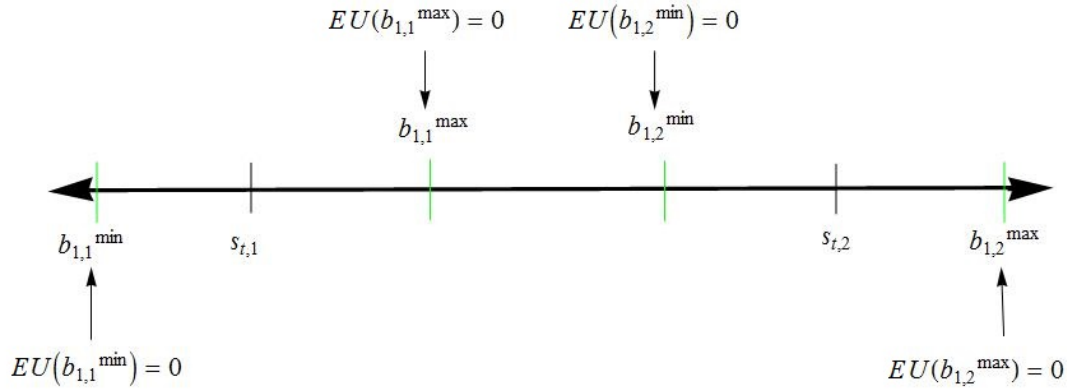
Under scenario 1 we can write the groups' payoff functions as:

$$\Pi_1 = \int_{b_1^{1,min}(\cdot)}^{b_1^{1,max}(\cdot)} \left(\pi_{1,1}(b_{1,i}) + \pi_{2,1}(b_{1,i}) \right) f_1(b_{1,i}) db_{1,i}$$

---

[13]Reports are determined by the truth, agendas and priors. The true state of the world and agendas remain unchanged over time. Groups should, therefore, not reverse roles and followers cannot switch whom they follow by assumption. In any way, there is no incentive in this model to switch allegiance.

[14]Note, that ω needs to be the same for both groups. The statement is, otherwise, not necessarily true, because groups would assign different weights in their slanting policies. Thus, I can only allow for groups to differ in their agendas, not in how they weigh up reader support and own agenda.
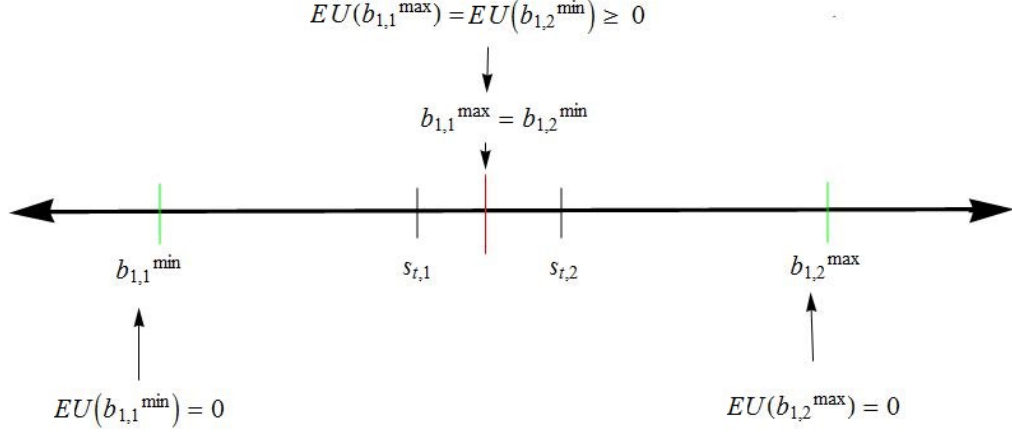
$$\Pi_2 = \int\limits_{b_1^{2,min}(\cdot)}^{b_1^{2,max}(\cdot)} \left(\pi_{1,2}(b_{1,i}) + \pi_{2,2}(b_{1,i}) \right) f_1(b_{1,i}) db_{1,i}.$$

$b_1^{j,min}(\cdot) = b_1^{j,min}(\boldsymbol{s_{t,j}(\tau)})$ and $b_1^{j,max}(\cdot) = b_1^{j,max}(\boldsymbol{s_{t,j}(\tau)})$ only depend on own slant and are independent of competitor slant. The payoff function is the same as under monopoly. Competition has, therefore, no effect on the optimal slanting policy of individual groups in this scenario.

It is sometimes argued that the existence of a left-leaning media group, such as MSNBC, encouraged the growth and rightward shift of Fox News and vice versa. I would, based on the above, argue that this might not be the case. These groups are so far apart that they were actually never competing for the middle ground. The shifts of these groups towards extremes over time might be better explained by proposition 2 than by increased competition. These groups might have had a fairly extreme agenda from the beginning. Their biases might become more extreme as they lead their followers towards their agenda slowly. The idea that competition is to blame for this development might be misleading.

Slanting policies $\boldsymbol{s_{t,1}(\tau)}$ and $\boldsymbol{s_{t,2}(\tau)}$ in the second scenario are close to each other. This is more likely to be the case when the agendas of the groups are fairly similar. There will be an agent i for whom $EU_{i,1} = EU_{i,2} \geq 0$; expected lifetime utility of following either group is identical and weakly positive. The groups are actively competing for the middle ground of types between them. Only two participation constraints are binding in this case and there exists an indifferent agent. This scenario is depicted in figure 4 below.

Figure 4: scenario 2

$$EU(b_{1,1}{}^{\max}) = EU\left(b_{1,2}{}^{\min}\right) \geq 0$$

$$b_{1,1}{}^{\max} = b_{1,2}{}^{\min}$$

$b_{1,1}{}^{\min}$ $s_{t,1}$ $s_{t,2}$ $b_{1,2}{}^{\max}$

$$EU\left(b_{1,1}{}^{\min}\right) = 0 \qquad\qquad EU(b_{1,2}{}^{\max}) = 0$$

The indifferent reader is characterised by $EU_{i,1}(\cdot) = EU_{i,2}(\cdot)$. Denote the type of the indifferent customer, which will depend on slanting policies of both media groups in both periods, as $b_1^I = b_1^I(s_{t,1}(\tau), s_{t,2}(\tau))$.

Hence, payoff functions can be written as:

$$\Pi_1 = \int_{b_1^{1,min}(\cdot)}^{b_1^I(\cdot)} \left(\omega EU(s_{t,1}(\tau), b_{1,i}) + (1-\omega)\Big(A_{1,1}(s_{1,1}(\tau)) + A_{2,1}(s_{2,1}(\tau))\Big) \right) f(b_{1,i}) db_{1,i}$$

$$\Pi_2 = \int_{b_1^I(\cdot)}^{b_1^{2,max}(\cdot)} \left(\omega EU(s_{t,2}(\tau), b_{1,i}) + (1-\omega)\Big(A_{1,2}(s_{1,2}(\tau)) + A_{2,2}(s_{2,2}(\tau))\Big) \right) f(b_{1,i}) db_{1,i}.$$

Maximisation yields the following four first order conditions (noticing again that expected lifetime utility of readers with a binding participation constraint is zero). The conditions, when solved, yield optimal duopolistic equilibrium slant $s_{t,j}(\tau)^{het,D}$.[15]

---

[15] $s_{t,j}(\tau)^{het,D}$ consists of $s_{t,1}(\tau)^{het,D}$ and $s_{t,2}(\tau)^{het,D}$, which in turn consist of $s_{1,1}(\tau)^{het,D}$ and $s_{2,1}(\tau)^{het,D}$, and $s_{1,2}(\tau)^{het,D}$ and $s_{2,2}(\tau)^{het,D}$ respectively. $s_{t,j}(\tau)$ also consists of the four equivalent individual slants.

$$\frac{\partial \Pi_1}{\partial s_{1,1}} = \int_{b_1^{1,min}(\boldsymbol{s_{t,1}(\tau)})}^{b_1^I(\boldsymbol{s_{t,j}(\tau)})} \left(\omega \frac{\partial EU_1(\boldsymbol{s_{t,1}(\tau)}, b_{1,i})}{\partial s_{1,1}} + (1-\omega) \frac{\partial A_{1,1}(s_{1,1}(\tau))}{\partial s_{1,1}} \right) f(b_{1,i}) db_{1,i}$$

$$+ \left(\pi_{1,1}(b_1^I(\boldsymbol{s_{t,j}(\tau)})) + \pi_{2,1}(b_1^I(\boldsymbol{s_{t,j}(\tau)})) \right) \left(f(b_1^I(\boldsymbol{s_{t,j}(\tau)})) \frac{\partial b_1^I(\boldsymbol{s_{t,j}(\tau)})}{\partial s_{1,1}} \right)$$

$$- \left(A_{1,1}(s_{1,1}(\tau)) + A_{2,1}(s_{2,1}(\tau)) \right) \left(f(b_1^{1,min}(\boldsymbol{s_{t,1}(\tau)})) \frac{\partial b_1^{1,min}(\boldsymbol{s_{t,1}(\tau)})}{\partial s_{1,1}} \right)$$

$$\frac{\partial \Pi_1}{\partial s_{2,1}} = \int_{b_1^{1,min}(\boldsymbol{s_{t,1}(\tau)})}^{b_1^I(\boldsymbol{s_{t,j}(\tau)})} \left(\omega \frac{\partial EU_1(\boldsymbol{s_{t,1}(\tau)}, b_{1,i})}{\partial s_{2,1}} + (1-\omega) \frac{\partial A_{2,1}(s_{2,1}(\tau))}{\partial s_{2,1}} \right) f(b_{1,i}) db_{1,i}$$

$$+ \left(\pi_{1,1}(b_1^I(\boldsymbol{s_{t,j}(\tau)})) + \pi_{2,1}(b_1^I(\boldsymbol{s_{t,j}(\tau)})) \right) \left(f(b_1^I(\boldsymbol{s_{t,j}(\tau)})) \frac{\partial b_1^I(\boldsymbol{s_{t,j}(\tau)})}{\partial s_{2,1}} \right)$$

$$- \left(A_{1,1}(s_{1,1}(\tau)) + A_{2,1}(s_{2,1}(\tau)) \right) \left(f(b_1^{1,min}(\boldsymbol{s_{t,1}(\tau)})) \frac{\partial b_1^{1,min}(\boldsymbol{s_{t,1}(\tau)})}{\partial s_{2,1}} \right)$$

$$\frac{\partial \Pi_2}{\partial s_{1,2}} = \int_{b_1^I(\boldsymbol{s_{t,j}(\tau)})}^{b_1^{2,max}(\boldsymbol{s_{t,2}(\tau)})} \left(\omega \frac{\partial EU_2(\boldsymbol{s_{t,2}(\tau)}, b_{1,i})}{\partial s_{1,2}} + (1-\omega) \frac{\partial A_{1,2}(s_{1,2}(\tau))}{\partial s_{1,2}} \right) f(b_{1,i}) db_{1,i}$$

$$+ \left(A_{1,2}(s_{1,2}(\tau)) + A_{2,2}(s_{2,2}(\tau)) \right) \left(f(b_1^{2,max}(\boldsymbol{s_{t,2}(\tau)})) \frac{\partial b_1^{2,max}(\boldsymbol{s_{t,2}(\tau)})}{\partial s_{1,2}} \right)$$

$$- \left(\pi_{1,2}(b_1^I(\boldsymbol{s_{t,j}(\tau)})) + \pi_{2,2}(b_1^I(\boldsymbol{s_{t,j}(\tau)})) \right) \left(f(b_1^I(\boldsymbol{s_{t,j}(\tau)})) \frac{\partial b_1^I(\boldsymbol{s_{t,j}(\tau)})}{\partial s_{1,2}} \right)$$

$$\frac{\partial \Pi_2}{\partial s_{2,2}} = \int_{b_1^I(\boldsymbol{s_{t,j}(\tau)})}^{b_1^{2,max}(\boldsymbol{s_{t,2}(\tau)})} \left(\omega \frac{\partial EU_2(\boldsymbol{s_{t,2}(\tau)}, b_{1,i})}{\partial s_{2,2}} + (1-\omega) \frac{\partial A_{2,2}(s_{2,2}(\tau))}{\partial s_{2,2}} \right) f(b_{1,i}) db_{1,i}$$

$$+ \left(A_{1,2}(s_{1,2}(\tau)) + A_{2,2}(s_{2,2}(\tau)) \right) \left(f(b_1^{2,max}(\boldsymbol{s_{t,2}(\tau)})) \frac{\partial b_1^{2,max}(\boldsymbol{s_{t,2}(\tau)})}{\partial s_{2,2}} \right)$$

$$- \left(\pi_{1,2}(b_1^I(\boldsymbol{s_{t,j}(\tau)})) + \pi_{2,2}(b_1^I(\boldsymbol{s_{t,j}(\tau)})) \right) \left(f(b_1^I(\boldsymbol{s_{t,j}(\tau)})) \frac{\partial b_1^I(\boldsymbol{s_{t,j}(\tau)})}{\partial s_{2,2}} \right)$$

Comparing these conditions with the ones faced by a monopolist with agenda a derived earlier, we can say that we are in scenario 1 when $b_1^{max}(\cdot)$ derived under monopoly is weakly smaller than $b_1^I(\cdot)$ (or $b_1^{min}(\cdot)$ is weakly larger for the group providing more positive slant). There is no such thing as an indifferent follower in that case. We know that competition will affect the group's optimal decision in all other cases. We also know that marginal benefits derived from either marginal reader must be equal. Otherwise, the group has an incentive to change its slanting policy and, therefore,

shift the marginal followers slightly. Lemmata 2 and 3 tell us that the group is going to slant according to the usual, linear slanting policy given the marginal followers.

Assume, in order to isolate from potential effects caused by the skewness of the distribution, that priors are uniformly distributed. We can see that three effects are going to determine whether the media groups will reduce or increase slant compared to the monopolistic case.

Firstly, groups face an incentive to capture the indifferent follower, because he is valuable as his expected utility is positive, and not zero, as is the case for the other marginal follower. This effect incentivises both groups to decrease slant and capture the indifferent follower.

Secondly, although the groups have an incentive to capture indifferent followers as they provide them with higher support, the groups tend to lose out to a larger followership at the other end of their marginal followers, i.e $\frac{\partial b_1^{1,min}(s_{t,1}(\tau))}{\partial s_{t,1}}$ (or for the second media group $\frac{\partial b_1^{2,max}(s_{t,2}(\tau))}{\partial s_{t,2}}$) tends to be larger than $\frac{\partial b_1^I(s_{t,j}(\tau))}{\partial s_{t,j}}$.[16] Where exactly these two effects cancel out depends on the parameters of the model.

Finally, one also has to consider that the existence of competition will change the average reader of groups. If $b_1^{2,max}$ and $b_1^{1,min}$ remained unchanged when compared to the monopolistic case, then the average reader of group one would now move to the left on the real line, while the average reader for group two would move to the right.[17]

We can, however, derive general results for a symmetric setting of agendas ($a_1 + a_2 = 0$) and a symmetric distribution of priors ($\bar{b}_1 = 0$). Suppose both groups always found it optimal to move towards the indifferent reader and, hence, end up sharing the market by positioning themselves at the same spot. They would both provide the identical slanting policies $S_t(\tau)$. The marginal followers for the first media group would be given by $f(b_1^{1,min}(S_t(\tau)))$ and $f(b_1^I(S_t(\tau))) = 0$. The marginal followers for the second group would be given by $f(b_1^I(S_t(\tau))) = 0$ and $f(b_1^{2,max}(S_t(\tau)))$. However, given these marginal followers, we know that the group's optimal slanting policy is not $S_t(\tau)$ but given by lemma 2 for the case of constant confirmation bias and lemma 3 for type-dependent confirmation bias. Group one will optimally slant towards own agenda a_1 with weight α_t^* and towards average reader prior $\frac{b_1^{1,min}(\cdot)}{2}$ with weight β_t^*. Group two will slant towards a_2 with weight α_t^* and average reader prior $\frac{b_1^{2,max}(\cdot)}{2}$ with weight β_t^*. It can, therefore, not be optimal for both groups to provide identical slant and be positioned at the centre of the distribution. Note

---

[16]I have so far not achieved to prove that this is generally the case for optimal slanting policies. However, it can easily be shown to hold in a symmetric setting (in which $a_1 + a_2 = 0$) for reasonable parameter values, i.e. $\bar{u}$ large enough to induce participation, $\chi > 0$ and $\phi > 0$.

[17]Once again, it is very hard to provide a general that this will necessarily always hold. It is theoretically conceivable that, for example, media group two decreases its slant so much that $b_1^{2,max}$ decreases so much that the average follower will actually move to the left, rather than to the right, on the real line.

that this holds even when both groups have no agenda at all. In that scenario they would provide zero bias in a monopolistic setting, but market segregation would drive them to bias the news when in competition. We can say that competition will increase slant provided by individual groups at least for media groups with modest agenda in a symmetric setting.

In summary, the existence of competition (assuming we are in scenario 2) shifts the average reader of each group further to the extremes, incentivising the groups to increase slant in competition. Capturing the indifferent reader might seem more valuable as he would provide positive support. However, increasing readership in that direction is harder than towards the extremes and costly because it would cause the media group to mismatch its new, more extreme audience. In any case, competition does not eliminate the supply side motive for slant in the case of heterogeneous followers. The following proposition summarises this analysis.

Proposition 7 (Duopolistic slant with heterogeneous readers).

- *Optimal duopolistic slanting policies are still of the linear form given by*

$$s_{t,j}(\tau)^{het,D} = \alpha_t^*(a_j - \tau) + \beta_t^*(b_1^{het,D} - \tau)$$

  if confirmation bias is type-independent.

- *Optimal duopolistic slanting policies are still of the linear form given by*

$$s_{t,j}(\tau)^{het,D} = \alpha_t^*(\hat{\phi}(s_{t,j}))(a_j - \tau) + \beta_t^*(\hat{\phi}(s_{t,j}))\left(b_1^{het,D} + \frac{c\hat{o}v_j(\phi, b_1)}{\hat{\phi}(s_{t,j})} - \tau\right)$$

  if confirmation bias is type-dependent.

- *Media groups will unambiguously increase their slant compared to the monopolistic benchmark in a symmetric setting $(a_1 + a_2 = 0)$ when $b_{1,i}$ is distributed symmetrically around $\bar{b}_1 = 0$.*

- *Competition has no effect on individual slanting policies when group's are sufficiently different.*

Slanting under competition in Mullainathan and Shleifer (2005) takes a very extreme form, because media groups compete in two dimensions. They compete in a basic Hotelling way by differentiating themselves in the biases they provide. They also compete over prices. Price competition is a lot fiercer in their model and groups will, thus, differentiate themselves with respect to their biases as much as possible, so that they can still charge high prices. Taking away the pricing mechanism makes the slant provided less extreme. However, I have shown that the supply side motive does not get eradicated and that the existence of competition provides incentives for media groups to increase biases. This effect is especially large when confirmation bias is type-dependent and stronger

for followers with more extreme views, because the market gets partitioned and groups focus on pleasing only their followers.

3 Conclusion

The literature focusing on supply side motives sees competition as unequivocally good. Mullainathan and Shleifer (2005) focus solely on the demand side motive and show that competition has no positive effects on biases, but a positive effect in the sense that it drives prices down. I provide a more nuanced view of this by combining both motives in a model without explicit pricing. In that model, we can see that competition still has an effect. Competition eliminates the supply side motive of slanting when readers have homogeneous beliefs. Whether the overall effect is positive or negative depends on the extremity of follower's prior compared to the groups' agendas.

I also provided a starting point for a more thorough analysis of dynamic effects. The analysis for the basic dynamics in the models showed that competition becomes more favourable over time, because the follower's prior converges towards the true state of the world. The same does not hold under monopolistic slant as long as the group's agenda is not aligned with the true state of the world.

If we consider it to be the group's role to provide factual reports and it is our aim for public opinion to be close to the true state of the world, then this could be, in the homogeneous setting, achieved by either competition or a monopolistic provider whose sole interest is the truth. Optimally, a monopolistic provider with high emphasis on the truth and an agenda that is aligned with the truth would be preferable, because its reports would be unbiased immediately rather than converging to the true state over time. This could be seen as the underlying logic and argument for the existence of large public media group. One could be sceptical about the practicality of such an unbiased monopolist. I have shown that, at least in the homogeneous case, one could argue for a competitive market, because it would not require individually unbiased media groups and ultimately lead to the same outcome.

A monopolist is not affected strongly by the heterogeneity of followers unless its agenda is strong and confirmation bias is type-dependent. Groups under competition react very differently compared to the homogeneous case, because the supply side motive is not eradicated. Different media groups will provide different slanting policies, because competition provides incentives for groups to differentiate themselves to increase their market shares. This will, in return, cause them to increase biases in order to satisfy their newly created, more extreme, audiences. I have shown that competition can,

thus, turn an unbiased monopolistic group into a biased duopolist. However, competition will have no effect at all when agendas are sufficiently different, because groups are not actively competing over the middle ground. The existence of the competitor has no influence on biases in that case.

In terms of the evolution of public opinion, we can say that the existence of competition will tend to engage the extremes more in the heterogeneous case, because readers who choose not to follow a group under monopoly will potentially choose to follow one under the more extreme biases provided under competition. This effect can be seen as positive, because these readers will adjust their beliefs over time and become less extreme. On the other hand, the middle ground of moderate beliefs could get eradicated slowly if groups actually compete. These readers with centrist views move to slightly more extreme views over time. Finally, the middle ground of modest priors remains abandoned under competition if groups do not actually compete.

The supply side motive does not get eliminated by competition when followers are heterogeneous. I would, therefore, argue that the case for a large public media group is more convincing in this case. However, it is vital that such a broadcaster puts a large emphasis on own agenda (small ω) and does not attempt to cater to the priors of followers. If the role of the media is to inform and educate, then trying to appeal to followers could potentially be harmful and lead to biases.

Appendices

Appendix A: Lemmata

We can use the following lemma, the equivalent to lemma 1 in Mullainathan and Shleifer (2005), defining agents' expected utilities, throughout the paper.

Lemma 1. *Let*

$$s_t(\tau) = \beta_t(B - \tau) + \alpha_t(A - \tau)$$

be the slanting policy of the media group.

The expected utility of a reader with prior b_1 is then given by:

$$EU(s_t(\tau), b_1) = \bar{u} - \chi\Big((\alpha_1 A + \beta_1 B)^2 + v_\tau(\alpha_1 + \beta_1)^2 + (\alpha_2 A + \beta_2 B)^2 + v_\tau(\alpha_2 + \beta_2)^2\Big)$$

$$- \phi\Big((\alpha_1 A + \beta_1 B - b_1)^2 + v_\tau(\alpha_1 + \beta_1 - 1)^2$$

$$+ (A(\alpha_2 - \gamma\alpha_1) + B(\beta_2 - \gamma\beta_1) - (1-\gamma)b_1)^2 + v_\tau(\alpha_2 + \beta_2 - \gamma(\alpha_1 + \beta_1) - (1-\gamma))^2\Big).$$

Hence:

$$\arg\max_{b_1} EU(s_t(\tau), b_1) = \frac{((1-\gamma(1-\gamma))\alpha_1 + (1-\gamma)\alpha_2)A + ((1-\gamma(1-\gamma))\beta_1 + (1-\gamma)\beta_2)B}{1 + (1-\gamma)^2}$$

and:

$$\tilde{b}_1 := \arg\max_{b_1} EU(s_t(\tau), b_1)\Big|_{B=b_1}$$

$$= -\frac{A\Big(\phi\alpha_2(\gamma - 1 - \gamma\beta_1 + \beta_2) + \alpha_1((\phi + \phi\gamma^2 + 2\chi)\beta_1 - \phi(1 - \gamma(1-\gamma) + \gamma\beta_2))\Big)}{(\phi + \phi\gamma^2 + 2\chi)\beta_1^2 - 2\phi\beta_1(1 - \gamma(1-\gamma) + \gamma\beta_2) + \phi(2 - 2\gamma + \gamma^2 - 2(1-\gamma)\beta_2 + \beta_2^2)}.$$

In addition, $EU(s_t(\tau), b_1) = 0$ defines two priors, $b_1^{min}(s_t(\tau))$ and $b_1^{max}(s_t(\tau))$ for which the participation constraint is binding. These priors are symmetric around the maximum of the expected utility function:

$$\frac{b_1^{min}(s_t(\tau)) + b_1^{max}(s_t(\tau))}{2} = \arg\max_{b_1} EU(s_t(\tau), b_1).$$

Proposition 1 shows that such a linear slanting policy is indeed optimal and solves for the specific α_t, β_t, A and B. The equation for $EU(s_t(\tau), b_1)$ shows the level of expected utility a follower with prior b_1 derives from this linear slanting policy.

$\arg\max_{b_1} EU(s_t(\tau), b_1)$ shows which follower will derive the largest utility under the proposed slanting policy. This is important to know, because this follower is the most valuable to the group as expected utility is transferred to the group in the form of support. The result differs only slightly from the one in Mullainathan and Shleifer (2005), because the views of followers develop over time and there exists a supply side motive. Thus, $\arg\max_{b_1} EU(s_t(\tau), b_1) \neq 0$ when $B = 0$ in my model while it would be zero in their model.

$\arg\max_{b_1} EU(s(\tau), b)\big|_{B=b_1}$ shows a similar thing. It shows which follower is most valuable (as he has the highest expected utility) conditional on the group using $B = b_1$ as its focal point, i.e. conditional on the group using a slanting policy that slants towards his prior. $\arg\max_{b_1} EU(s(\tau), b)\big|_{B=b} = 0$ in the single period setting of Mullainathan and Shleifer (2005). This is because there is no supply side motive in their model ($\alpha_1 = \alpha_2 = 0$).[18] In general, the group has an incentive to match a follower's prior, because utility is maximised for the follower whose prior is matched by slant (this

---

[18]Furthermore, $\gamma = \beta_2 = 0$ in Mullainathan and Shleifer (2005).

is due to the confirmation bias). However, expected utility is larger the closer a reader's prior is to the expected true value of the world ($E(\tau) = 0$), because followers dislike slant. This means that groups have a stronger incentive to match priors close to the true state of the world.

The follower whose prior, taking into account its evolution, is best matched by the slanting policy over both periods has the highest expected utility. However, the group would also prefer the matched follower to have a modest prior. The prior for which expected utility is maximised when matched is defined by $\tilde{b}_1$. This follower is the most valuable to groups in terms of potential support. With $A = 0$, i.e. if we render agenda unimportant, $\tilde{b}_1 = 0$ as in Mullainathan and Shleifer (2005).

The final statement of the lemma makes use of the symmetric property of the expected utility function and shows that the participation constraints are binding at two priors. These are equidistant from the maximum of the expected utility function.

Proof. $EU(\boldsymbol{s_t}(\boldsymbol{\tau}), b_1)$ is derived by plugging the proposed linear slanting policy into the expected utility of a reader followed by basic simplification (using $E(\tau) = 0$ and $E(\tau^2) = v_\tau$ when solving the integral). It is then used to derive to general results for linear slanting policies. $\arg\max_{b_1} EU(\boldsymbol{s_t}(\boldsymbol{\tau}), b_1)$ is derived by straightforward maximisation. $\tilde{b}_1$ is again derived by simple maximisation subject to $B = b_1$ in $\boldsymbol{s_t}(\boldsymbol{\tau})$. The final statement simply applies the fact that the expected utility function, which is in essence a basic loss function, is symmetric around its maximum. $\square$

The following lemma, which builds on lemma 2 in Mullainathan and Shleifer (2005), will help us, in combination with lemma 1, understand the nature of the optimal slanting policy of a monopolistic groups when followers are heterogeneous.

Lemma 2. *Let b_1^{max} and b_1^{min} be the biases of two readers, with $b_1^{min} < b_1^{max}$. If all readers with $b_1^{min} \leq b_{1,i} \leq b_1^{max}$ follow the group, then the strategy*

$$\boldsymbol{s_t}(\boldsymbol{\tau}) = \boldsymbol{\alpha_t}^*(a - \tau) + \boldsymbol{\beta_t}^*(\hat{b}_1 - \tau)$$

maximises the group's objective function, where $\hat{b}_1 = \frac{\int_{b_1^{min}}^{b_1^{max}} b_{1,i} f_1(b_{1,i}) db_{1,i}}{F(b_1^{max}) - F(b_1^{min})}$ is the average prior of followers.

The lemma shows that the group will use the usual linear slanting policy if we hold the marginal followers fixed. Whatever the marginal readers are, it will assign the same weights as in the

homogeneous case and will use the average follower prior as one of its focal points (with own agenda being the other).

Proof. The group's payoff function is given by:

$$\int_{b_1^{min}}^{b_1^{max}} \Big(\omega EU(s_t(\tau), b_{1,i}) + (1 - \omega) \left(A_1(s_1(\tau)) + A_2(s_2(\tau)) \right) \Big) f(b_{1,i}) db_{1,i}.$$

Taking derivatives with respect to $s_1(\tau)$ and $s_2(\tau)$ results in the following first order conditions after some basic rearrangements:

$$s_1 = \frac{1 - \omega}{1 - \omega(1 - \phi - \chi - \phi\gamma^2)}(a - \tau) + \frac{\omega\phi}{1 - \omega(1 - \phi - \chi - \phi\gamma^2)}(\gamma s_2 + (1 - \gamma(1 - \gamma))(\hat{b}_1 - \tau))$$

$$s_2 = \frac{1 - \omega}{1 - \omega(1 - \phi - \chi)}(a - \tau) + \frac{\omega\phi}{1 - \omega(1 - \phi - \chi)}(\gamma s_1 + (1 - \gamma)(\hat{b}_1 - \tau))$$

which, when solved, yields the lemma. $\qquad\square$

Lemma 3. *Let b_1^{max} and b_1^{min} be the biases of two readers, with $b_1^{min} < b_1^{max}$. If all readers with $b_1^{min} \leq b_{1,i} \leq b_1^{max}$ follow the group, then the strategy*

$$s_t(\tau) = \alpha_t^*(\hat{\phi}(s_t))(a - \tau) + \beta_t^*(\hat{\phi}(s_t)) \left(\hat{b}_1 + \frac{c\hat{o}v(\phi, b_1)}{\hat{\phi}(s_t)} - \tau \right)$$

maximises the groups objective function Π, where $\hat{\phi}(s_t) = \dfrac{\int_{b_1^{min}}^{b_1^{max}} \phi(b_{1,i}) f_1(b_{1,i}) db_{1,i}}{F(b_1^{max}) - F(b_1^{min})}$ denotes average confirmation bias of all followers for which the participation constraint holds and $c\hat{o}v(\phi, b_1)$ denotes the covariance between the confirmation bias and followers' views, but only for followers of type $b_1^{max} \geq b_{1,i} \geq b_1^{min}$.

Proof of lemma 3. This lemma extends lemma 2 to the case of type-dependent confirmation bias. Taking derivatives of the objective with respect to $s_1(\tau)$ and $s_2(\tau)$ results in the following first order conditions after some basic rearrangements:

$$s_1 = \frac{1 - \omega}{1 - \omega(1 - \hat{\phi}(s_t) - \chi - \phi\gamma^2)}(a - \tau)$$

$$+ \frac{\omega\hat{\phi}(s_t)}{1 - \omega(1 - \hat{\phi}(s_t) - \chi - \hat{\phi}(s_t)\gamma^2)} \left(\gamma s_2 + (1 - \gamma(1 - \gamma)) \left(\frac{\int_{b_1^{min}}^{b_1^{max}} \phi(b_{1,i}) b_{1,i} f_1(b_{1,i}) db_{1,i}}{F(b_1^{max}) - F(b_1^{min})} - \tau \right) \right)$$

$$s_2 = \frac{1-\omega}{1-\omega(1-\hat{\phi}(\boldsymbol{s_t})-\chi)}(a-\tau)$$

$$+ \frac{\omega\hat{\phi}(\boldsymbol{s_t})}{1-\omega(1-\hat{\phi}(\boldsymbol{s_t})-\chi)}\left(\gamma s_1 + (1-\gamma)\left(\frac{\int_{b_1^{min}}^{b_1^{max}}\phi(b_{1,i})b_{1,i}f_1(b_{1,i})db_{1,i}}{F(b_1^{max})-F(b_1^{min})}-\tau\right)\right).$$

Solving these and simplifying by using the definition for the covariance, i.e. $cov(x,y) = E(xy) - E(x)E(y)$, results in the lemma. $\qquad\square$

Appendix B: Proof of propositions

Proof of proposition 1. Maximising the simplified objective function

$$\Pi_{hom} = \omega EU_{hom} + (1-\omega)\int_\tau \left(\bar{a} - (s_1+\tau-a)^2 - (s_2+\tau-a)^2\right)h(\tau)d\tau$$

with respect to the group's choices s_1 and s_2, using the updating rule $b_2 = \gamma(s_1+\tau)+(1-\gamma)b_1$ and noticing that there are no interdependencies and, hence, maximisation can be pointwise, yields:

$$\frac{\partial\Pi_{hom}}{\partial s_1} = -2\omega\chi s_1 - 2\omega\phi(s_1+\tau-b_1) - 2(1-\omega)(s_1+\tau-a) + 2\omega\phi\gamma(s_2+\tau-b_1) = 0$$

$$\frac{\partial\Pi_{hom}}{\partial s_2} = -2\omega\chi s_2 - 2\omega\phi(s_2+\tau-b_2) - 2(1-\omega)(s_2+\tau-a) = 0.$$

Solving the system of two equations yields the results summarised in the proposition.

The follower's participation constraint is stricter by assumption. Hence, the equilibrium outlined in proposition 1 exists as long as $EU(\boldsymbol{s_t^*}(\boldsymbol{\tau}),b_1)$, derived by substituting $\boldsymbol{s_t^*}(\boldsymbol{\tau})$ for $\boldsymbol{s_t}(\boldsymbol{\tau})$ in lemma 1, is weakly positive. The equilibrium defined in proposition 1 exists if $b_1^{min}(\boldsymbol{s_t^*}(\boldsymbol{\tau})) \leq b_1 \leq b_1^{max}(\boldsymbol{s_t^*}(\boldsymbol{\tau}))$.[19] Intuitively, the equilibrium will be the solution to an unconstrained maximisation problem for the media group when the follower's prior is not extremely different from own agenda and/or the agent does not care about the truth. When the follower's view is very extreme and he puts a high emphasis on hearing news that match his prior, then the media group might have to provide him with news that match his opinion sufficiently to induce participation. However, if this extreme prior is opposite to own agenda, inducing participation comes at a cost to the media group. $\qquad\square$

---

[19]The group will, otherwise, provide a slanting policy that makes the follower's participation constraint binding by shifting its policy further towards the agent's prior and/or the expected truth and away from its own agenda to increase the agent's expected utility and, thus, induce participation. The optimal slanting policy would, hence, be a constrained optimum for the group.

Proof of proposition 2. Since combined weights go down over time, $\alpha_1^* + \beta_1^* \geq \alpha_2^* + \beta_2^*$, there is a higher implicit weight on the truth in later periods. In a symmetric setting ($|a - \tau| = |b_1 - \tau|$) slant is, therefore, decreasing over time.

The weight on matching the follower's prior goes down over time, while the weight on own agenda is increasing, $\beta_1^* \geq \beta_2^*$ and $\alpha_1^* \leq \alpha_2^*$. This means that bias can only increase over time when $|a - \tau|$ is sufficiently larger than $|b_1 - \tau|$. When an increasing weight is placed on the motive for bias that creates more extreme bias, the this will tend to increase overall bias.

We can derive the actual set for which slant is increasing over time. As shown in proposition 1, the weights in the optimal slanting policy change over time, while the points around which the group slants remain the same. We can write the change in slant over time as:

$$s_2^*(\tau) - s_1^*(\tau) = \Delta s^*(\tau) = (\beta_2^* - \beta_1^*)(b_1 - \tau) + (\alpha_2^* - \alpha_1^*)(a - \tau) = \Delta \beta^*(b_1 - \tau) + \Delta \alpha^*(a - \tau).$$

From this we get that

$$\Delta s^*(\tau) \geq 0 \quad \text{if } (a - \tau) \geq -\frac{\Delta \beta^*}{\Delta \alpha^*}(b_1 - \tau) = \frac{1 - \omega + \chi\omega}{1 - \omega}(b_1 - \tau)$$

which tells us how slant changes over time. It grows when agenda does not offset the follower's prior and is sufficiently more extreme than it. What really matters, though, is the absolute value of slant and its change ($\Delta |s^*(\tau)|$) which will depend on whether slant was positive or negative to start with.

Using the same procedure as above, we can derive a condition for $s_1^*(\tau) \geq 0$ and $s_2^*(\tau) \geq 0$ which are given by:

$$s_1^*(\tau) \geq 0 \quad \text{if } (a - \tau) \geq -\frac{\beta_1^*}{\alpha_1^*}(b_1 - \tau)$$

$$s_2^*(\tau) \geq 0 \quad \text{if } (a - \tau) \geq -\frac{\beta_2^*}{\alpha_2^*}(b_1 - \tau).$$

Therefore, we know that slant is unambiguously increasing over time when the condition for $s_1^*(\tau) \geq 0$ and $\Delta s^*(\tau) \geq 0$ are satisfied. Slant must also increase when $\Delta s^*(\tau) \geq 0$ is satisfied, while both $s_1^*(\tau) \geq 0$ and $s_2^*(\tau) \geq 0$ are violated. What happens when slant switches the sign between periods is not generally clear. We know, though, that absolute slant must always go up if $s_1^* = s_2^*$ and $\Delta s^*(\tau) \geq 0$ is satisfied. It proves difficult to analyse the behaviour of the condition $s_1^* = s_2^*$ directly. However, we know that the line at which $s_1^* = s_2^*$ must lie between the lines depicting $s_1^* = 0$ and $s_2^* = 0$, which we have derived above. We also know that the line depicting $s_1^* = 0$ is steeper than the one depicting $s_2^* =$, because $-\frac{\beta_1^*}{\alpha_1^*} < -\frac{\beta_2^*}{\alpha_2^*}$.

Furthermore, we know that:

$$\frac{\partial\left(-\frac{\Delta\beta^*}{\Delta\alpha^*}\right)}{\partial\chi}, \frac{\partial\left(-\frac{\Delta\beta^*}{\Delta\alpha^*}\right)}{\partial\omega} \geq 0, \frac{\partial\left(-\frac{\Delta\beta^*}{\Delta\alpha^*}\right)}{\partial\phi}, \frac{\partial\left(-\frac{\Delta\beta^*}{\Delta\alpha^*}\right)}{\partial\gamma} = 0, \frac{\partial\left(-\frac{\beta_1^*}{\alpha_1^*}\right)}{\partial\chi}, \frac{\partial\left(-\frac{\beta_1^*}{\alpha_1^*}\right)}{\partial\phi}, \frac{\partial\left(-\frac{\beta_1^*}{\alpha_1^*}\right)}{\partial\omega} \leq 0.$$

These imply that the line depicting the set of agendas and priors for which $s_1^* = s_2^*$ becomes steeper with ϕ, ω and χ. Since $\frac{\partial\frac{1-\omega+\chi\omega}{1-\omega}}{\partial\phi} = 0$, $\frac{\partial\frac{1-\omega+\chi\omega}{1-\omega}}{\partial\omega} > 0$ and $\frac{\partial\frac{1-\omega+\chi\omega}{1-\omega}}{\partial\chi} > 0$, we know that the condition on $\Delta s^*(\tau) = 0$ also becomes steeper (its slope is positive). The final statement of the proposition follows, because the area between the two conditions is the set of agendas and priors for which slant decreases in absolute terms $\qquad\square$

Proof of proposition 3. At $t = 1$, expected lifetime utility from following either media/lobby group has to be equal:

$$EU_{1,hom}(\boldsymbol{s_{t,1}(\tau)}) = EU_{2,hom}(\boldsymbol{s_{t,2}(\tau)}) = EU_{hom}.$$

It also has to be maximised, because the competitor could otherwise attract the entire market by providing a slanting policy that gives the reader higher utility. Maximisation yields the following first order conditions:

$$\frac{\partial EU_{hom}}{\partial s_1} = -2\chi s_1 - 2\phi(s_1 + \tau - b_1) - 2\phi(s_2 + \tau - b_2)(-\frac{\partial b_2}{\partial s_1}) = 0$$

$$\frac{\partial EU_{hom}}{\partial s_2} = -2\chi s_2 - 2\phi(s_2 + \tau - b_2) = 0.$$

Notice that the assumption of loyalty does not affect the results here, because $\frac{\partial Eu(s_2)}{\partial s_2} = \frac{\partial EU}{\partial s_2}$. We could also have said that agents must be provided the same expected utility at $t = 1$ and at $t = 2$, rather than the same expected lifetime utility. It would yield the same first order conditions. Simplification leads to:

$$s_1 = \frac{\phi}{\phi + \chi + \phi\gamma^2}(\gamma s_2 + (1 - \gamma(1 - \gamma))(b_1 - \tau))$$

$$s_2 = \frac{\phi}{\phi + \chi}(\gamma s_1 + (1 - \gamma)(b_1 - \tau)).$$

The first condition is identical to the optimal slanting strategy of a media group in Mullainathan and Shleifer (2005) in the homogeneous case with $\gamma = 0$, i.e. it matches the static example of a media group without own agenda and no evolution of public opinion.

The proposition follows from solving the two first order conditions. The condition for the existence of the equilibrium in proposition 3, i.e. that the market gets served, is given by $EU(\boldsymbol{s_t^D(\tau)}, b_1)$, which can be derived by substituting $\boldsymbol{s_t^D(\tau)}$ for $\boldsymbol{s_t(\tau)}$ in lemma 1, being weakly positive. Since the follower's expected lifetime utility is maximised under competition, the equilibrium will exist if $\overline{u}$

is large enough for $EU(s_t^D(\tau), b_1)$ to be weakly positive. Otherwise, the market will not be served. If the market is not served under competition, it will not be served under any market structure.

The final statement of the proposition follows directly from $\beta_1^D \geq \beta_2^D$. We can also see that the weight on the follower's prior is larger under competition as $\beta_t^D \geq \beta_t^*$. $\qquad\square$

Proof of proposition 4. Write the difference in slant in monopoly versus slant under competition in period t as:

$$s_t^*(\tau) - s_t^D(\tau) = \beta_t^*(b_1 - \tau) + \alpha_t^*(a - \tau) - \beta_t^D(b_1 - \tau)$$

After rearranging we can find a condition on the parameters for slant to be smaller in monopoly than in competition:

$$s_t^*(\tau) - s_t^D(\tau) \leq 0 \quad \text{if } b_1 \leq -\frac{\alpha_t^*}{\beta_t^* - \beta_t^D}(a - \tau) + \tau.$$

We can also rearrange $s_t^*(\tau)$ and identify a condition for which monopolistic slant is positive in t:

$$s_t^*(\tau) \geq 0 \quad \text{if } b_1 \geq -\frac{\alpha_t^*}{\beta_t^*}(a - \tau) + \tau.$$

Finally, we know from proposition 3 that:

$$s_t^D(\tau) \geq 0 \quad \text{if } b_1 \geq \tau.$$

If all three condtions are satisfied simultaneously, then slant is clearly smaller in absolute terms under monopoly. Since all weights are positive, the condition $b_1 \geq \tau$ (slant under competition is positive) is clearly stricter than the condition $b_1 \geq -\frac{\alpha_t^*}{\beta_t^*}(a - \tau) + \tau$ (slant under monopoly is positive) for $a > \tau$. $b_1 \geq \tau$ and $b_1 \leq -\frac{\alpha_t^*}{\beta_t^* - \beta_t^D}(a - \tau) + \tau$ can be satisfied simultaneously for $a > \tau$ as long as $\beta_t^* < \beta_t^D$ which I have shown to be true. This shows that there exist some b_1 for any $a > \tau$ for which slant under monopoly is smaller in absolute terms. Filling in the weights for the periods, one can see that the set of b_1 that satisfy this gets smaller. Graphically speaking, all lines depicting the conditions above become flatter. This is driven by the fact that the weight on agenda is increasing over time in monopoly, i.e. slant provided does not converge to zero, while the supply side motive is eliminated in competition.

The same approach can be applied to the case of $a < \tau$ or to the case in which $s_t^D(\tau) \leq 0$, $s_t^*(\tau) \leq 0$ and $s_t^*(\tau) - s_t^D(\tau) \geq 0$. I will not do all these scenarios in this proof as the steps would be identical to the ones shown above and I have already established the existence of a set of parameters for which the statement in the proposition holds. $\qquad\square$

Proof of proposition 5. Lemma 2 shows that the linear slanting policy with $A = a$ and $B = \hat{b}_1$ and the same weights as in the homogeneous case $\boldsymbol{\alpha_t^*}$ and $\boldsymbol{\beta_t^*}$ maximises the monopolist's payoff for fixed marginal readers b_1^{min} and b_1^{max}. Given the marginal readers, i.e. disregarding the effect slanting has on the participation constraints, the monopolistic group behaves like in the homogeneous case. It slants towards average reader prior.

However, the types of marginal readers depend on the slanting policy itself ($b_1^{min} = b_1^{min}(\boldsymbol{s_t(\tau)})$ and $b_1^{max} = b_1^{max}(\boldsymbol{s_t(\tau)})$) in the heterogeneous case and the group has to take that into account. Lemma 2 shows that the group uses the linear slanting policy outlined in the lemma whatever the marginal readers. It slants towards $\hat{b}_1$ with weight β_t^* and towards own agenda a with weight $\boldsymbol{\alpha_t^*}$. It remains unclear, though, what $\hat{b}_1$ is. We can, however, see that optimal slant will equalise the benefits of marginal readers, i.e. the second lines of the first order conditions are equal to zero. Given this, the media group can then use the slanting policy outlined in the lemma 2 to maximise its benefits.

If $b_{1,i}$ is symmetrically distributed around $\tilde{b}_1$, we know that the average prior $\bar{b}_1$ is equal to $\tilde{b}_1$ defined in lemma 1. Since by lemma 1, $b_1^{min}(\cdot)$ and $b_1^{max}(\cdot)$ are symmetrically placed around $\tilde{b}_1$ it follows that the average prior of readers following the group $\hat{b}_1$ is therefore also equal to the overall average $\bar{b}_1$. The group finds it optimal to slant around this point, because by lemma 2 it maximises expected reader utility and by lemma 1 it maximizes expected utility. As shown in lemma 2 for given marginal readers the group wants to slant towards $\hat{b}_1$. However, by lemma 1 expected utility is rising for priors being closer to $\tilde{b}_1$. In this scenario the two coincide though.

If $\bar{b}_1 = \tilde{b}_1$ but the distribution is not symmetric this does not hold, because $\hat{b}_1 \neq \bar{b}_1$ even though possibly $\hat{b}_1 = \tilde{b}_1$, i.e. the average prior of those following the group does not coincide with the unconditional average prior of all readers. In this scenario the group will shift b_1^{het} further towards the marginal reader that is represented stronger in the population, because it provides the group with an opportunity to spread its agenda to a wider audience.[20]

If $\bar{b}_1 \neq \tilde{b}_1$, then for given marginal readers as shown in lemma 2 the group would slant towards $\hat{b}_1$. However, maximised expected utility by lemma 1 is higher when slanting towards a point closer to $\tilde{b}_1$, because these readers are closer to the expected true state of the world. Given the symmetric shape of the expected utility function, a higher maximum implicitly also means a wider audience, i.e. $b_1^{max} - b_1^{min}$ is larger.[21] This means that by shifting the slanting point closer towards $\tilde{b}_1$ the

---

[20]This can also be seen in the first-order conditions. In this scenario $\frac{\partial b_1^{max}(s_t(\tau))}{\partial s_t} = \frac{\partial b_1^{min}(s_t(\tau))}{\partial s_t}$, but $f(b_1^{max}(s_t(\tau))) \neq f(b_1^{min}(s_t(\tau)))$.

[21]This can also be seen in the first-order conditions. In this scenario $\frac{\partial b_1^{max}(s_t(\tau))}{\partial s_t} \neq \frac{\partial b_1^{min}(s_t(\tau))}{\partial s_t}$.

group can increase its payoff because readers' utilities increase and the audience widens. The group has to trade this off against mismatching its audience. This is why overall $b_1^{het} = \hat{b}_1$ is determined by $\bar{b}_1$ and $\tilde{b}_1$.

This allows us to summarise the optimal monopolistic slanting policy when faced with heterogeneous readers as done in the proposition.

$\square$

Proof of proposition 6. Lemma 3 shows that the optimal slanting policy is still linear, but will now take the relation between priors and confirmation bias, as well as the fact that the average confirmation bias of readers depends on the slanting policy itself, into account. The weights $\boldsymbol{\alpha_t^*}$ and $\boldsymbol{\beta_t^*}$ are still similar to the homogeneous case, but are not simple constants determined by the parameters of the model. The confirmation bias of followers depends on priors of the followers for whom the participation constraint holds, which in turn depends on the slanting policy. The weights are, therefore, functions of slanting policy themselves.

The overall effect of non-constant confirmation bias is not clear cut. However, we can derive some general tendencies and results for special cases. Suppose, the monopolist would just adopt s_t^{het} (remember that $\hat{b}_1 = b_1^{het}$) as its policy and not consider that confirmation bias is type-dependent. The covariance term has no influence on the optimal decision when $\phi(\cdot)$ is symmetric around $s_t(\tau)^{het}$. Although confirmation bias depends on type, it does not incentivise the media group to change its decision, because the effect is symmetric and affects marginal benefits both at the upper and lower end of types in the same fashion. If, furthermore $\hat{\phi}(s_t(\tau)^{het}) = \phi$, i.e. the conditional confirmation bias matches the constant confirmation bias, the weights are also unaffected. It follows that $s_t(\tau)^{het}$ would still be the optimal slanting policy in this case.

Evidence from the psychology literature suggests that confirmation bias is increasing in the extremity of own belief. So, assume that $\phi(b_{1,i})$ is symmetrically u-shaped around $b_{1,i} = 0$. This means that, if the media groups slant was modest, the logic just explained would apply and the influence of type-dependent confirmation bias would be negligible.

In summary, for $\phi(b_{1,i})$ being symmetrically u-shaped around $b_{1,i} = 0$, media groups with positive slant face a positive covariance for their respective readers and media groups with a negative slant a negative covariance. This effect provides an incentive to increase absolute slant further when confirmation bias is type-dependent, because readers with extreme views require to have their

priors matched disproportionately. However, we also have to consider the effect that the average confirmation bias of readers depends on the slanting policy itself, which affects the weights $\boldsymbol{\alpha_t^*}$ and $\boldsymbol{\beta_t^*}$.

The first effect causes the group with $s_t(\tau)^{het} > 0$ to further increase its slant because of a positive covariance between types and confirmation bias. As shown earlier:

$$\frac{\partial \beta_1^*}{\partial \phi} \geq 0, \frac{\partial \alpha_1^*}{\partial \phi} \leq 0; \quad \frac{\partial \beta_2^*}{\partial \phi} \geq 0, \frac{\partial \alpha_2^*}{\partial \phi} \leq 0$$

$$\left| \frac{\partial \beta_1^*}{\partial \phi} \right| \geq \left| \frac{\partial \alpha_1^*}{\partial \phi} \right|; \quad \left| \frac{\partial \beta_2^*}{\partial \phi} \right| \geq \left| \frac{\partial \alpha_2^*}{\partial \phi} \right|.$$

This means that the weight on matching priors $(\boldsymbol{\beta_t^*})$ increases, while the weight on own agenda $(\boldsymbol{\alpha_t^*})$ decreases when confirmation bias is increasing in extremity of priors, $\left. \frac{\partial \hat{\phi}(s_t)}{\partial s_t} \right|_{s_t > 0} > 0$ and $\left. \frac{\partial \hat{\phi}(s_t)}{\partial s_t} \right|_{s_t < 0} < 0$. Whether this effect causes slant to further increase, thus, depends on whether agenda or reader priors are more extreme.[22] The effect will cause the group to increase bias even further if followers' priors are more extreme than own agenda. $\qquad \square$

References

Anderson, Simon P.; and McLaren, John. "Media Mergers and Media Bias with Rational Consumers", *Journal of the European Economic Association*, 2012, 10(4), pp. 831-859

Ashworth, Scott; and Shotts, Kenneth W. "Does informative media commentary reduce politicians' incentive to pander?", *Journal of Public Economics*, 2010, 94(11), pp. 838-847

Austen-Smith, David. "Campaign Contributions and Access", *American Political Science Review*, 1995, 89(3), pp. 566-581

Baron, David P. "Persistent media bias", *Journal of Public Economics*, 2006, 90, pp. 1-30

Bennedsen, Morten; and Feldmann, Sven E. "Lobbying Legislatures", *Journal of Political Economy*, 2002, 110(4), pp. 919-946

Bernhardt, Dan; Krasa, John; and Polborn, Mattias. "Political polarization and the electoral effects of media bias", *Journal of Public Economics*, 2008, 92, pp. 1092-1104

---

[22]The example works just the other way around for a group faced with a negative covariance between confirmation bias and priors.

Besley, Timothy; and Coate, Stephen. "Lobbying and Welfare in a Representative Democracy", *Review of Economic Studies*, 2001, 68(1), pp. 67-82

Chan, Jimmy; and Suen, Wing. "A Spatial Theory of News Consumption and Electoral Competition", *Review of Economic Studies*, 2008, 75, pp. 699-728

Gentzkow, Matthew; and Shapiro, Jesse M. "Media Bias and Reputation", *Journal of Political Economy*, 2006, 114(2), pp. 280-316

Gentzkow, Matthew; and Shapiro, Jesse M. "What drives media slant? Evidence from US daily newspapers", *Econometrica*, 2010, 78(1), pp. 35-71

Gerber, Alan S.; Karlan, Dean; and Bergan, Daniel. "Does the Media Matter? A Field Experiment Measuring the Effect of Newspapers on Voting Behavior and Political Opinions", *American Economic Journal: Applied Economics*, 2009, 1(2), pp. 35-52

Jones, Martin; and Sugden, Robert. "Positive confirmation bias in the acquisition of information", *Theory and Decision*, 2001, 50(1), pp. 59-99

Klayman, Josh. "Varieties of confirmation bias", *Psychology of learning and motivation*, 1995, 32, pp. 385-418

Kuhn, Deanna; and Lao, Joseph. "Effects of evidence on attitudes: Is polarization the norm?", *Psychological Science*, 1996, pp. 115-120

Li, Ming; and Mylovanov, Tymofiy. "Indifferent Public, Passionate Advocates, and Strategic Media", *unpublished paper*, 2007

Lord, Charles G.; Ross, Lee; and Lepper, Mark A. "Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence", *Journal of Personality and Social Psychology*, 1979, 37(11), pp. 2098-2109

Maskin, Eric; and Tirole, Jean. "The politician and the judge: Accountability in government", *American Economic Review*, 2004, pp. 1034-1054

Mullainathan, Sendhil; and Shleifer, Andrei. "The Market for News", *American Economic Review*, 2005, 95(4), pp. 1031-1053

Mynatt, Clifford R.; Doherty, Michael E.; and Tweney, Ryan D. "Confirmation bias in a simulated research environment: An experimental study of scientific inference", *The quarterly journal of experimental psychology*, 1977, 29(1), pp. 85-95

Pew Research Center for the People and the Press. "Views of the News Media: 1985-2011", *http://www.people-press.org/files/legacy-pdf/9-22-2011%20Media%20Attitudes%20Release.pdf*

Sobbrio, Francesco. "Indirect Lobbying and Media Bias", *Quarterly Journal of Political Science*, 2011, 6(3-4), pp. 235-274

Zaller, John (ed.). "The nature and origins of mass opinion", *Cambridge university press*, 1992

The Effect of Group Size and Composition on Mechanisms with Voluntary Contributions

Kimon Doulis

School of Economics

University of Edinburgh

Abstract

The existing literature analysing voluntary contribution mechanisms for public goods has shown that individual payments are decreasing in the size of the population, but that aggregate payments are increasing. This paper uses a simple, discrete model with a binary public good and focuses on individual and expected aggregate payments. I derive results in line with the literature for symmetric mechanisms, in which agents are identical ex ante. A thorough analysis is provided for a basic two agent case and results are then generalised to an n agent environment. A specific focus is put on explaining the ex ante budget constraint a designer faces and on how it affects his choice of mechanism, especially when varying the level of provision costs. Efficiency can be measured in several ways in this model. The paper focuses mainly on explaining aggregate payments, because these will affect all measures positively.

The paper then compares these results with the outcomes of an asymmetric mechanism, in which agents are heterogeneous at the ex ante stage. The random variable, which contains the private information, is not distributed according to the same distribution function for all agents. Some agents are perceived to be more likely to value the public good highly. A form of price discrimination can be used in such a situation. The paper aims to provide a basic insight into how this might affect the choice of mechanism. I present results in a basic two agent case and show how making use of price discrimination can potentially enhance efficiency. The scope to use price discrimination is shown to be limited in large groups and to disappear asymptotically. Nevertheless, heterogeneous groups tend to outperform their homogeneous equivalents in larger groups.

1   Introduction & Literature

The existence of the free rider problem has been well established in economic theory, in laboratory experiments and can be observed in real life. Individuals face an incentive to contribute less than socially optimal whenever they have to contribute monetarily or with effort towards the provision of a good that benefits an entire group.

To overcome this problem, which potentially causes under-provision of public goods, a typical solution is to implement a first best allocation by not considering individual rationality or participation

constraints. Contributions towards the public good or specific public projects are collected via involuntary taxation. Such mechanisms are efficient and relatively simple to implement. However, the power to coerce agents into making payments does not always exist.[1] These mechanisms can also create distributionary issues, because some agents potentially derive negative benefits when forced to make a payment.[2]

Voluntary contribution mechanisms

Analysing voluntary contribution mechanisms requires agents to report their types to the designer. Research establishing the revelation principle, such as Myerson (1979), enables us to focus on direct mechanisms. It is, therefore, sufficient to look for payment and allocation rules that satisfy participation, incentive compatibility and budget constraints.

D'Aspremont and Gérard-Varet (1979) were among the first to show the existence of such a mechanism for a public good with a maximum provision level of one and characterise such a mechanism. They establish that the mechanism could solve the collective decision problem efficiently if participation constraints are only of the ex ante type. Ex post efficiency can be implemented with a mechanism satisfying incentive compatibility as long as participation constraints are only ex ante. This model is extended by Ledyard and Palfrey (1999) who allow for non-identical weights for different agents in the social welfare functions. The designer cares intrinsically about who pays how much. Ledyard and Palfrey (1999) characterise the mechanism for this broader setting and derive similar results to d'Aspremont and Gérard-Varet (1979). Since the first best solution depends on the weights assigned to different agents, the optimal way to reduce the burden for low-valuation agents is implemented by reducing the provision level of the public good.[3]

While early papers like d'Aspremont and Gérard-Varet (1979) emphasise the existence and potential efficiency of voluntary contribution mechanisms when participation constraints are only ex ante, the appearance of the Myerson-Satterthwaite theorem (Myerson and Satterthwaite, 1983; Myerson, 1984) caused a shift in the literature of mechanism design and asymmetric information bargaining. Simply put, it shows, for a private good, that there exists no incentive compatible, individually rational mechanism satisfying a budget constraint that ensures ex post efficient trade between a buyer and a seller, if valuations are private information. This inefficiency result caused authors to also revisit the case for public goods.

---

[1] I do not have the coercive power to force my flatmates into contributing towards a new TV that could benefit everyone in the flat. I also do not have the power to coerce people into doing the dishes, although clean dishes would benefit everyone.

[2] Maybe my flatmates actually like dirty dishes and do not value a TV.

[3] A maximum provision level of one in both these papers is reasonably interpreted as a binary public good with risk-neutral agents, where intermediate provision levels just stand for the probability of provision.

Group size and composition

Parts of the individual's efforts or contributions towards the public good cause only external benefits, which individuals account for insufficiently. This suggests that the free rider problem could be increasing in the group size. If a group project has to be handed in in groups of two, then each member of the group has an incentive to slack off and benefit from the other person's work. However, individual effort has a large effect on the group's performance. The individual's perceived influence over the outcome is, therefore, large. If, in contrast, the group consists of 10 people, the individual's perceived influence over the eventual group project seems much more limited. This tends to increase the free rider problem in large groups. A similar logic holds for the case of voluntary contributions towards public projects, which are studied in this paper.

The economic literature about voluntary contribution mechanisms for public goods has largely focused on very complex characterisations of incentive compatible and individually rational contribution mechanisms and derived several asymptotic (in)efficiency results. As the group size increases and the ability of individual agents to influence the provision level of the public good disappears, individual payments that agents are willing to make decrease. The literature then splits and derives different results from this. One part of the literature focuses on mentioning that aggregate payments could still increase and that, if the public good is provided once aggregate payments reach a specific threshold, voluntary contribution mechanisms can achieve efficiency, as long as the group is large enough. Others focus on the fact that the gap in contributions between a first best solution and the voluntary contribution mechanism is widening with group size and that voluntary contribution mechanisms become more inefficient with group size.

Mailath and Postlewaite (1990) establish a result that can be seen as the public good equivalent of the Myerson-Satterthwaite theorem. They show that there exists no incentive compatible, individually rational mechanism satisfying a budget constraint ensuring ex post efficiency even with just two agents in a bargaining setting with asymmetric information about valuations for a continuous public good. Furthermore, they show that the probability of implementing an efficient allocation using a voluntary mechanism is decreasing in the group size and converges towards zero. The paper establishes a fundamental, negative result similar to Myerson and Satterthwaite (1983). Not only are voluntary contribution mechanisms inefficient, but the level of inefficiency is increasing in group size as well. Individual contributions are decreasing in group size, because the level of pivotality is negatively associated with group size.

However, Hellwig (2003) points out that Mailath and Postlewaite (1990) assume a provision cost for the public good that depends on the number of participants. While this might be reasonable

in some scenarios, it does not have to be. Hellwig (2003) also shows that aggregate payments can be increasing, although individual payments are decreasing due to a decrease in individual influence or pivotality. The level of inefficiency depends strongly on assumptions made about costs and on the kind of public good considered. Hellwig (2003), assuming a constant provision cost independent of n, moves on to show that the asymptotic inefficiency result of Mailath and Postlewaite (1990) holds for continuous public goods, but not for a threshold or binary public good. Although total contributions in a voluntary contribution mechanism as a fraction of the first best level are decreasing in group size, they are still increasing in total terms. When it is only necessary to reach a specific threshold of contributions in order to provide the good, this threshold will be eventually reached by voluntary contributions. Hellwig (2003) also considers the case of excludable public goods, as does Norman (2004) in a similar model, and establishes that the threat of exclusion always makes agents pivotal with regards to whether they get access to the good or not. The ability to exclude agents from access to the public good can, thus, lead to a significant increase in the efficiency level of voluntary contribution mechanisms.

Al-Najjar and Smorodinsky (2000) establish bounds on the pivotality of agents depending on group sizes in a broader setting, which can also be used to analyse voting behaviour. Their results can be used to provide bounds for the cost function under which Hellwig's results for a threshold public good hold. I am, like Hellwig (2003), going to assume that provision costs are independent of group size. The effect of the functional form of the cost function has been well established. The main focus in this paper is going to be on simplifying the existing literature, and on analysing the effect of ex ante heterogeneous agents.

Ex ante heterogeneity

The existing literature has also largely focused on agents who are heterogeneous in their valuations for the public good, but are not ex ante heterogeneous.[4] Ex ante heterogeneity means that agents' valuations are random draws out of potentially different distribution functions. Some agents look to others to be more likely to value the good or to be more likely to have a high valuation than others.

This distinction appears important, because it is typically not the case that individual members of a group all appear to be the same, even if some information is held privately and unknown to other members of the group. Some of my flatmates have very clean rooms, some do not. Some of them have large collections of DVDs in their rooms, others do not. Hence, although I do not know how much they value clean dishes or a new TV, it is simply not the case that I have the same belief over

---

[4]Some papers do consider ex ante heterogeneity and I will specifically point them out in this section.

their valuations. The same can hold for large groups. When deciding whether to build a bridge for cars connecting two towns, the planner can observe who owns a car and who does not. Beliefs over other agents' valuations might, thus, differ according to some observable characteristic. The main aim of this paper is to analyse how ex ante heterogeneity affects the free-riding problem and contributions towards public projects. They should be affected, because the perceived influence or pivotality of agents over the provision of the public good depends on their beliefs over other agents' valuations. This analysis aims to provide a step towards answering the question of whether the existence of ex ante heterogeneity within groups is beneficial for the efficiency level of voluntary contribution mechanisms or not.

While Mailath and Postlewaite (1990) allow for the valuations to be from different distribution functions and, hence, allow for ex ante heterogeneity, it does not matter to their asymptotic results. Total contributions of ex ante heterogeneous groups as a fraction of their first best counterpart still converge to zero. Hellwig (2003) as well as Al-Najjar and Smorodinsky (2000) do not allow for ex ante heterogeneity. While ex ante heterogeneity might not matter asymptotically in Mailath and Postlewaite (1990), because of their assumption about the cost function, it should influence the level of contributions for given group sizes. Furthermore, the speed of convergence should also depend on the order of agents. This paper aims to provide an argument for this and provides some quantitative results, highlighting the effect of ex ante heterogeneity. In general, the paper shows that more costly projects can be financed in heterogeneous groups, or that projects with identical costs are more likely to be provided in heterogeneous groups.

Experimental evidence

There exists a wider experimental literature asking similar questions and the work presented here can provide some theoretical foundation to it. One of the earlier attempts to analyse empirically how voluntary contributions depend on own wealth and the expected behaviour of other agents was conducted by Rapoport and Suleiman (1993). In a basic public good contribution game, their experimental results suggest that agents with a larger initial endowment contribute more in total terms. Each group in the experiment consists of five members with different endowments. Participants know that the assignment of endowments is random and know all the endowments they could get. They know whether they are relatively poor or rich after receiving their endowment. It remains unclear in the results whether individuals contribute more or less relatively to their endowment as the endowment changes. Furthermore, their results suggest that the individual's estimate of average relative contributions of the rest of the group is independent of the own endowment, which implies that the estimate for total contributions by the rest of the groups are decreasing in

the own endowment. Although this setting does not fully capture the idea of ex ante heterogeneous agents presented in this paper, and wealth and valuation for the public good are not the same, the experiment is useful, because it shows that estimates about other agents' contributions affect own contributions negatively. In a simple theoretical framework, I am going to propose that it is not the difference in endowments that drives the differences in contributions, but the different estimates about other agents' valuations and different levels of perceived pivotality.

Isaac, Walker and Williams (1994) conduct an experiment studying a basic public good provision game in which agents' endowments are homogeneous. They focus on the effect group size might have on contribution levels and on the efficiency level of public good provision. In general, their results suggest that total contributions are increasing in group size and that efficiency levels are, therefore, increasing as well.[5] It is important to highlight that the first only implies the second if the provision costs are independent of group size. These results partially support Hellwig's (2003) as well as Al-Najjar and Smorodinsky's (2000) argument that aggregate payments should increase, even though individual free-riding incentives might increase with group size.

Fisher et al. (1995) study the case in which group members in small groups face different returns on their investments in the public good. They compare contributions to the case in which all members of the group face identical returns on their investments. The results show that average contributions in the heterogeneous groups, in which half the group faces a high and half a low return, lie between the average contributions of groups in which all members face a high or low return. Unsurprisingly, agents facing a higher return within a heterogeneous group contribute more, but less than they do when in a homogeneous group in which everyone faces a high return.

The case of different returns on investments is comparable to the theoretical setting of ex ante heterogeneous groups presented here. Fundamentally, both settings imply that own contributions seem more important or valuable for some agents. However, the comparisons in Fisher et al. (1995) between heterogeneous and homogeneous groups do not appear to be completely fair. In order to compare the two in terms of efficiency, one should not compare how the allocations in each case compare as a fraction to first best allocations in each case respectively, but one should compare a heterogeneous group to a homogeneous group for which the first best allocation is identical. Only then can one make a fair judgment on which group setting tends to do better. This is the kind of comparison that I aim to provide in a theoretical framework.

Chan et al. (1999) provide further insight into how heterogeneity affects contributions. They

---

[5]Their data suggests partially that even individual contributions are increasing in group size. This is very hard to reconcile with any model of utility maximising agents, because the free rider problem should be increasing at the individual level.

find that average contributions are larger in groups in which endowments and preferences are heterogeneous. An advantage to Fisher et al. (1995) is that Chan et al. (1999) try to make sure that heterogeneous groups are compared to equivalent homogeneous groups. However, they restrict themselves to very small groups (n=3) and only report group contributions. We cannot say whether the group contributions are increasing in heterogeneous groups because all members increase their contributions, or whether members with high preferences for the public good increase their contributions disproportionately. Cherry et al. (2005) report different results by showing that heterogeneity in wealth endowments on its own decreases average contributions. The theoretical claim presented here will be that groups with larger heterogeneity in preferences, not necessarily in wealth, will tend to be better at providing threshold public goods.

The experiments presented above consider a basic public good contribution game, in which individuals can contribute some of their private endowment into a public purse, which yields a specific return. In this paper, I am going to focus on a threshold or binary public good. A public project needs a specific level of funding and is only provided if that threshold level of funding is met. There is no continuous level of the public good, it is either provided or not. Contributions towards such a threshold public good have been experimentally analysed by Cadsby and Maynes (1999), be it in a setting of full information and homogeneous preferences and endowments. Incomplete information and heterogeneity in preferences are crucial to the analysis and Cadsby and Maynes (1999) is, therefore, only partially relevant. The results are, nevertheless, still relevant, because they show that agents behave as expected for threshold public goods. Furthermore, Cadsby and Maynes (1999) show that results under basic voluntary contribution mechanisms can be improved upon significantly if a money-back guarantee is implemented in the case the good is not provided. Rondeau et al. (1999) provide similar results and show that allocations can be further improved upon if a proportional rebate of excess contributions is implemented as well.

Finally, one should note that almost all public good contribution experiments show that individuals in labs contribute more than theory would suggest. This indicates that some agents are motivated by more than simply their pivotality. As shown by Fischbacher et al. (2001), some agents match the rational agent model while others can be classified as conditional contributors, a behaviour that can be explained by reciprocity. Other motives that could explain larger contributions than typically predicted in theory are altruism or warm-glow. Interestingly, Burlando and Guala (2005) present experimental evidence suggesting that the existence of different types of agents is detrimental to public good contributions. They show that contributions in heterogeneous groups consisting of all types (co-operators, reciprocal agents etc.) are smaller than in homogeneous groups. However, this could be due to a learning effect in their experimental setup. Furthermore, the theoretical framework

presented here will focus on rational agents without any form of social preferences. These might exist and, as shown, affect results in the experimental literature, but my main focus in this paper lies in the analysis of voluntary contribution mechanisms without relying on asymptotic results and comparing ex ante homogeneous to ex ante heterogeneous groups. Rational agent models without social preferences still provide a valuable baseline model for this, especially if additional motives such as altruism are independent of group size or composition.

Overall contributions

The existing literature focuses on proving the existence of an individually rational, incentive compatible contribution mechanism in different scenarios, characterising it and describing the asymptotic behaviour of the allocation implemented by it.

This paper aims to contribute to the literature by:

- explaining existing results of the literature in a more intuitive way by presenting more detailed and simplified derivations;
- presenting a simplified model that allows for the analysis of efficiency levels for given group sizes at different provision costs;
- analysing the effect of ex ante heterogeneity on efficiency levels in small and large groups; and
- providing a comparison of efficiency levels for ex ante heterogeneous and comparable homogeneous groups.

2 The Model

There are n agents bargaining about the provision of a public good Q, with $Q \in \{0, 1\}$. In other words, the decision about the provision of the public good is a binary, yes-or-no, decision. The cost of providing the public good is given by $c(Q)$ and is independent of the number of agents, i.e. $c(1) = c$ and $c(0) = 0$.

Agents are risk-neutral and their utility function is given by $u(Q) = Q$. They differ in their valuation for the public good, which is given by θ_i for agent i and is private information. Agent i derives a net benefit, $U(Q)$, of:

$$U(Q) = \theta_i u(Q) - P_i$$

where P_i is his contribution or payment. For simplification, I will assume that θ_i can only take two values. There are only agents with high or low valuations.

In the *homogeneous setting*, the valuation θ_i is the realisation of a random variable $\tilde{\theta}_i$, which is distributed according to a distribution function f and a cumulative distribution function F. Distribution functions are identical for all agents. Specifically, I will assume that individual valuations are Bernoulli distributed and that the low valuation is zero. This means that $f_i(\theta_i)=f(\theta_i)$ is given by:

$$\theta_i = \left\{ \begin{array}{ll} \theta_H & \text{with } \pi \\ 0 & \text{with } 1 - \pi \end{array} \right. .$$

In the homogeneous setting, individual valuations are either high (θ_H) with probability π, or low/zero with probability $1 - \pi$.

The population is divided into two subpopulations in the *heterogeneous setting*. Subpopulation A consists of n_A members and subpopulation B of n_B members. The total population size is given by $n_A + n_B = n$. The size of each group and whether an agent is a member of subpopulation A or B is known to all agents. Individual valuations within each subpopulation are, as in the homogeneous setting, distributed according to the same distribution function. However, the distribution functions differ across groups.

For all members of subpopulation A, $f_A(\theta_i)$ is given by:

$$\theta_i = \left\{ \begin{array}{ll} \theta_H & \text{with } \pi_A \\ 0 & \text{with } 1 - \pi_A \end{array} \right. .$$

For all members of subpopulation B, $f_B(\theta_i)$ is given by:

$$\theta_i = \left\{ \begin{array}{ll} \theta_H & \text{with } \pi_B \\ 0 & \text{with } 1 - \pi_B \end{array} \right. .$$

Without loss of generality, members of subpopulation A are assumed to have higher average valuations, $\pi_A > \pi_B$. There is no ex post difference between a member of A or B when they both have high or both have low valuations. However, there is a difference ex ante. A member of subpopulation A looks more likely to others to have a high valuation.

Any mechanism implemented specifies two things:

- A *payment rule*. In this paper, interim payments are going to depend only on own reported valuations $P_i(\theta_i)$. Setting payments depending on all types is considerably more effort. The designer would have to revisit agents, because payments are contingent on all types. Intuitively, we rule out revisiting agents and it is, therefore, not possible to set payments contingent on other agents' valuations. If communication between agents is impossible or too costly, one might also interpret this as agents not trusting the designer's motives and his reports of other agents' types.[6]

_____

[6]One could set payments depending an all reported valuations $p_i(\theta_i, \boldsymbol{\theta_{-i}})$, where $\boldsymbol{\theta_{-i}}$ is the vector of other valuations. However, due to the risk neutrality of agents, what matters to the individual agent in the interim

- An *allocation rule*. The rule is of the following kind in this model: $Q = 1$ if $\#\theta_H \geq n - l$ and $Q = 0$ otherwise. This means that the good gets provided if a sufficient number of agents (more or equal to $n-l$ of them) report to be high types and it does not get provided otherwise. The interim payments, $P_i(\theta_i)$, will depend on the provision rule, which are defined by l, as well as group size and probabilities. We can, thus, write $P_i(\theta_i) = P_i(n, l, \pi, \theta_i)$.

It is the designer's objective to set a mechanism that maximises the provision probability for a given level of provision costs c. We will see that the provision probability is increasing in l. In effect, the designer will aim to implement the mechanism with the highest l possible subject to feasibility constraints. The level of c will, in this model, determine which mechanisms are feasible and which are not.

2.1 Homogeneous agents

This part of the paper will derive results for the homogeneous case. It introduces the key equations, such as participation, incentive compatibility and budget constraints in detail to establish the key principles involved. Since payments, in the homogeneous setting, must be identical for agents with identical valuations, the payment function further simplifies to $P_i(\cdot) = P(\cdot)$.

In this binary simplification of Mailath and Postlewaite (1990) and Hellwig (2003) the mechanism is interim **individually rational** (satisfies participation constraints) if:

$$E\left[\theta_i u(Q(\theta_i, \boldsymbol{\theta_{-i}})) - P(n, l, \pi, \theta_i)\right] \geq 0$$

where expectations are taken over $\boldsymbol{\theta_{-i}}$ (the vector of valuations of other agents).

The mechanism is interim **incentive compatible** if agents prefer to reveal their true type θ_i rather than pretending to be of type $\hat{\theta}_i$, i.e. if:

$$P(n, l, \pi, \theta_i) - P(n, l, \pi, \hat{\theta}_i) \leq E\left[\theta_i u(Q(\theta_i, \boldsymbol{\theta_{-i}}) - \theta_i u(Q(\hat{\theta}_i, \boldsymbol{\theta_{-i}})\right].$$

One can, in theory, introduce different budget constraints for feasibility of the mechanism. The mechanism would be feasible and satisfy the **ex post budget constraint** if:

$$\sum_{i=1}^{n} P(n, l, \pi, \theta_i) \geq c(Q(\boldsymbol{\theta}))$$

for all $\boldsymbol{\theta}$, where $\boldsymbol{\theta}$ is the n-dimensional vector of valuations. Aggregate payments have to cover provision costs ex post in all cases.

One can introduce a weaker **ex ante budget constraint** meaning that the mechanism is feasible

incentive compatibility constraint is the expected payment. $P_i(\theta_i)$ is just the expectation of $p_i(\theta_i, \boldsymbol{\theta_{-i}})$. I set a single payment $P_i(\theta_i)$, because it satisfies the interim incentive compatibility constraints for the same values.

if:

$$E\left[\sum_{i=1}^{n} P(n, l, \pi, \theta_i)\right] \geq E\left[c(Q(\boldsymbol{\theta}))\right]$$

which requires payments to cover costs only in expectation. Individual payments are not refundable when the good is not provided and payments above the level necessary for provision are also not refunded (no partial refund) when it is. We can, thus, write $EAP(n, l, \pi) = E\left[\sum_{i=1}^{n} P(n, l, \pi, \theta_i)\right]$ as expected aggregate payments and write the ex ante budget constraint as:

$$EAP(n, l, \pi) \geq E\left[c(Q(\boldsymbol{\theta}))\right].$$

Using a specific mechanism, which, as explained, is characterised by a cost-splitting rule l, will determine individual payments $P(n, l, \pi, \theta_i)$, which in term determine expected aggregate payments $EAP(n, l, \pi)$. The designer will sometimes provide the good but get insufficient funding. Sometimes funding will be higher than necessary. On average, costs will be covered. Agents will sometimes make payments which are lost and not refunded. Sometimes they will make payments and derive a positive utility. The payment they will be willing to make is then determined by the incentive compatibility and individual rationality constraints. Throughout the paper, I will usually, for simplicity, use the ex ante budget constraint unless mentioned otherwise.[7]

Because $c(1) = c$ and $Q = 1$ for $\#\theta_H \geq n - l$ and zero otherwise, as explained in the allocation rule, we can write the simplified **budget constraint** as:

$$EAP(n, l, \pi) \geq c \times \Pr(\#\theta_H \geq n - l)$$

or

$$c \leq \frac{EAP(n, l, \pi)}{\Pr(\#\theta_H \geq n - l)} = \bar{c}(n, l, \pi)$$

where $\Pr(\#\theta_H \geq n - l)$ is the provision probability for the public good. Expected aggregate payments implemented by the allocation rule (cost-splitting rule), summarised by l, and individual payments $P(n, l, \pi, \theta_i)$, which we still have to derive, are determined by the incentive compatibility and individual rationality constraints.[8]

---

[7]So do Mailath and Postlewaite (1990) and Hellwig (2003). Cramton, Gibbons and Klemperer (1987) show for a continuous type-set that the two budget constraints are equivalent when participation constraints are interim. For any allocation specifying a payment structure and an allocation rule under the ex ante budget constraint there exists a payment structure implementing the same allocation rule ex post. Essentially, the two are equivalent if the designer has access to fair insurance. Boergers and Norman (2009) generalise this result for very general settings. Even without the existence of a third party insurer, the existence of risk neutral agents means that the agents themselves can provide fair insurance yielding the equivalence result.

We can see that the individual rationality and incentive compatibility constraints depend on interim payments $P(\theta_i)$, which are essentially average ex post payments. Intuitively, one can set strictly positive ex post payments, $p(\theta_i, \boldsymbol{\theta_{-i}})$, even for low types in the case many other agents turn out to be low types as well. One can compensate agents by reducing ex post payments if the other agents are high types as well. This works as long as $P(\theta_i)$ remains unchanged. This provides insurance for the case that fewer than expected agents turn out to value the good ex post. Note, that even risk-neutral low types with zero valuation would be willing to provide insurance as long as the expected ex post payment, or $P(\theta_L)$, satisfies the incentive compatibility and individual rationality constraints.

[8]These equations are crucial in understanding the choice of mechanism, the asymptotic behaviour of voluntary

We can see clearly that the provision probability for the public good is strictly increasing in l, meaning that the optimal mechanism (the one that maximises provision probability) is simply the one that requires the smallest amount of people to have high valuations. However, such a mechanism will only create individual payments determined by the incentive compatibility constraints. It is, therefore, only feasible if it satisfies the budget constraint above. We will establish that expected aggregate payments are decreasing in l up to some point, which depends on n and π. It follows that the provision probability maximising mechanism might not be feasible. It is optimal to always choose a cost-splitting rule that requires as few individual payments for provision as possible (an l as large as possible), but still satisfies the budget constraint.

Holding l constant, the argument by Mailath and Postlewaite (1990) essentially is that the costs are a function of n, $c(n)$. Although $EAP(n, l, \pi)$ might increase in n, the costs are increasing faster, making the budget constraint stricter. This leads to the asymptotic result that no mechanism can provide the good with positive probability. Hellwig (2003) points out that all that really matters asymptotically is the behaviour of $EAP(n, l, \pi)$ compared to the behaviour of $\Pr(\cdot)$ when costs are independent of group size. He shows both of these to be increasing in n. However, aggregate payments are increasing faster. It follows that the budget constraint becomes less strict for a given c. The provision probability will, therefore, reach one in large enough groups. The group size required to ensure provision is, however, much larger than under first best. In any case, understanding how expected payments depend on group size and composition is crucial as any measure of efficiency (provision probabilities holding costs constant, or the level of costs holding the probability constant) depends on it. Besides deriving a full set of results for the two agent case, it is the aim of this paper to compare the behaviour of EAP in the homogeneous case with the case of heterogeneous agents.

I provide a simpler approach to the contrasting results of Hellwig (2003) and Mailath and Postlewaite (1990) by tracking EAP. Using the simplified setting provided here, I provide an idea for the (in)efficiency levels in finite group sizes. I also compare results derived under homogeneity with results derived in a setting of ex ante heterogeneity, which Hellwig (2003) does not allow for.[9]

A simple **ex ante measure of efficiency** would be to provide the good as long as expected aggregate valuations outweigh provision costs, $E\left[\sum_{i=1}^{n} \theta_i\right] \geq c$, or:

$$Q = 1 \quad \text{iff } c \in \left[0, E\sum_{i=1}^{n} \theta_i\right].$$

A simple **ex post measure of efficiency** would be to provide the good if actual aggregate valuations outweigh costs, $\sum_{i=1}^{n} \theta_i \geq c$.

Using the ex ante measure, this gives us the following first best provision probabilities from an ex ante perspective:

$$\Pr(Q = 1) = \begin{cases} \Pr(\#\theta_H \geq 1) & \text{if } c \in [0, \theta_H] \\ \Pr(\#\theta_H \geq 2) & \text{if } c \in [\theta_H, 2\theta_H] \\ \vdots & \vdots \\ \Pr(\#\theta_H \geq n) & \text{if } c \in [(n-1)\theta_H, n\theta_H] \\ 0 & \text{if } c > n\theta_H \end{cases}.$$

In general, efficiency could be defined in several measures for small groups. This paper will have a look at a variety of them.[10] Some voluntary contribution mechanisms might yield a high provision probability, but only exist for low c. The opposite might be true for others. One measure of efficiency could be the provision probability given a level of costs. Another could be the highest level of costs for which a voluntary contribution mechanism yielding positive provision probability exists. As shown in the simplified budget constraint, all measures are related positively to EAP. I will, therefore, focus on tracking EAP while reporting provision probabilities in small groups. The analysis will show that the two measures are not only related, but asymptotically identical and that efficiency is fully summed up in EAP.

2.1.1 Two agents

I am going to provide a detailed step-by step derivation for this most basic of cases. The case with n agents and heterogeneous agents build on the same fundamental logic. Following the derivations in detail will provide an intuitive understanding of the results for the more complicated cases. The two agent example highlights in an intuitive fashion the causes of inefficiencies of voluntary contribution mechanisms.

Taking expectations over the other agent's type, we can rewrite the general incentive compatibility and individual rationality constraints derived earlier as shown below.

Incentive compatibility constraint for agent i, if $\theta_i = \theta_H$:

$$P(\theta_H, \cdot) - P(\theta_L, \cdot) \leq \theta_H \bigg(\pi u(Q(\theta_H, \theta_H)) + (1 - \pi)u(Q(\theta_H, \theta_L)) - \pi u(Q(\theta_L, \theta_H)) - (1 - \pi)u(Q(\theta_L, \theta_L)) \bigg).$$

Individual rationality constraint for agent i, if $\theta_i = \theta_H$:

$$P(\theta_H, \cdot) \leq \theta_H \bigg(\pi u(Q(\theta_H, \theta_H)) + (1 - \pi)u(Q(\theta_H, \theta_L)) \bigg).$$

---

[10]There might be a mechanism with a low provision probability that satisfies the budget constraint for high values of c. Is this more or less efficient than a mechanism with high provision probability that only exists for low c?

Incentive compatibility constraint for agent i if $\theta_i = \theta_L$:

$$P(\theta_L, \cdot) - P(\theta_H, \cdot) \le \theta_L \Big(\pi u(Q(\theta_L, \theta_H)) + (1-\pi)u(Q(\theta_L, \theta_L)) - \pi u(Q(\theta_H, \theta_H)) - (1-\pi)u(Q(\theta_H, \theta_L)) \Big).$$

Individual rationality constraint for agent i if $\theta_i = \theta_L$:

$$P(\theta_L, \cdot) \le \theta_L \Big(\pi u(Q(\theta_L, \theta_H)) + (1-\pi)u(Q(\theta_L, \theta_L)) \Big).$$

The second term in the brackets in $u(Q(\theta_L, \theta_H))$ is the other agent's valuation, i.e. θ_j. The constraints for agent j are symmetrical in the homogeneous case.

With our assumption that low valuation agents derive zero utility from the public good, $\theta_L = 0$, it is obvious that the payments for a low valuation agent must be given by:

$$P_L = P(\theta_L, \cdot) = 0$$

independent of the other agent's type. Clearly, the inequality from the individual rationality constraint has to be binding. It is not the case that $P_L < 0$, because such a transfer would affect the budget negatively and also make it more difficult for the incentive compatibility constraint of high types to hold; both of these affecting efficiency levels negatively. As usual in these settings, the participation constraint of low types and the incentive compatibility constraint of high types are the binding constraints, because high types have an incentive to mimic low types. We can focus solely on payments of high types in our simplified, binary setting.

Using the notation $P_H(\cdot) = P(\theta_H, \cdot)$, we can use $P_L = 0$ to rewrite the **incentive compatibility constraint of the high type** as:

$$P_H(\cdot) \le \theta_H \Big(\pi u(Q(\theta_H, \theta_H)) + (1-\pi)u(Q(\theta_H, \theta_L)) - \pi u(Q(\theta_L, \theta_H)) - (1-\pi)u(Q(\theta_L, \theta_L)) \Big)$$

and the **individual rationality constraint** as:

$$P_H(\cdot) \le \theta_H \Big(\pi u(Q(\theta_H, \theta_H)) + (1-\pi)u(Q(\theta_H, \theta_L)) \Big).$$

Since utility is non-negative, we can see that the incentive compatibility is stricter and implies the participation constraint. Using our utility function $u(Q) = Q$ in combination with risk neutrality, we can think of utility as the provision probability.

This allows us to rewrite the **incentive compatibility** and **individual rationality constraints** as:

$$P_H(\cdot) \le \theta_H \Big(\pi \big(Pr(Q=1|\theta_H, \theta_H) - Pr(Q=1|\theta_L, \theta_H) \big) + (1-\pi) \big(Pr(Q=1|\theta_H, \theta_L) - Pr(Q=1|\theta_L, \theta_L) \big) \Big)$$

and

$$P_H(\cdot) \le \theta_H \Big(\pi \big(Pr(Q=1|\theta_H, \theta_H) \big) + (1-\pi) \big(Pr(Q=1|\theta_H, \theta_L) \big) \Big).$$

Note that the provision of Q is deterministic given a vector $\boldsymbol{\theta}$, which here is just equal to θ_j.

Expectations and probabilities are taken over θ_j. This makes clearer why the incentive compatibility constraint is stricter. The reason is that the public good might be provided even if the high type falsely reports to be a low type if $\theta_j = \theta_H$.[11] Since any payments contribute positively to the budget, the incentive compatibility constraint will hold with equality and payments of high types $(P_H(\cdot) = P_H(n, l, \pi))$ are, thus, given by the incentive compatibility constraint above. In general, as outlined earlier, the solution to this equation will depend on the provision rule, determined by l, as this will determine the provision probabilities (0 or 1 for each of the cases).

In the *first mechanism*, the interim payment of high types, $P_H(n, l, \pi)$, is specified to be so large that a single contribution will cover the costs of the binary public good, i.e. $P_H(n, l\pi) \geq c$. It follows from the incentive compatibility constraint that the high type will not make a payment ex post if the other agent has reported to be a high type. This is the case, because j's payment will cover the costs and i will, thus, prefer to free-ride. However, if the other agent reports a low type, then the agent is pivotal for the provision of the public good. The incentive compatible payment determined by this mechanism, for which $l = 1$, as only $n - l$ payments are necessary for provision, is then given by:

$$P_H(2, 1, \pi) = \theta_H \left(\pi \left(1 - 1 \right) + (1 - \pi) \left(1 - 0 \right) \right) = (1 - \pi)\theta_H.$$

Such a mechanism produces ex ante aggregate payments of:

$$EAP(2, 1, \pi) = 2\pi P_H = 2\pi(1 - \pi)\theta_H$$

and a provision probability of:

$$\Pr(\#\theta_H \geq 1) = \pi^2 + 2\pi(1 - \pi).$$

The mechanism is only feasible if it satisfies the budget constraint, $c \leq \bar{c}(n, l, \pi)$. However, the mechanism also has to satisfy the initial condition, $c \leq P_H$. This condition is stricter in this instance, because $(1 - \pi)\theta_H < \bar{c}(2, 1, \pi)$. The mechanism in which a single payment is sufficient exists only as long as costs are below the upper threshold of $\bar{c}_1$:

$$c \leq \bar{c}_1 = (1 - \pi)\theta_H.$$

In the *second mechanism*, the interim payment of high types is specified to be such that it is necessary for both agents to contribute in order to provide the good, i.e. $P_H < c \leq 2P_H$ and $l = 0$. Costs are split. Applying the same logic as previously, agent i is only pivotal if agent j reports to be a high type as well. The public good is not provided when j reports to be a low type. In that case i will not make a payment, because any payment he makes is not refundable and strictly

---

[11]Whether that is the case or not depends on the payments for high types. If a single payment is sufficient to cover the costs, then the goods will still be provided.

utility reducing. The incentive compatible payment is:

$$P_H(2, 0, \pi) = \pi\theta_H.$$

This mechanism produces ex ante aggregate payments of:

$$EAP(2, 0, \pi) = 2\pi^2\theta_H$$

and a provision probability of:

$$\Pr(\#\theta_H \geq 2) = \pi^2.$$

However, such an interim incentive compatible and individual rational mechanism, in which costs are split, only satisfies the ex ante budget constraint for:

$$c \leq \bar{c}(2, 0, \pi) = 2\theta_H.$$

It is one of the aims of this paper to compare mechanisms for given group sizes. Let us *compare the two mechanisms* explained above in this basic two agent setting. It is optimal to set a mechanism requiring only a single payment for $c \in [0, \bar{c}_1]$, because the provision probability is larger for that mechanism. Hence, an optimal mechanism yields a provision probability of $\pi^2 + 2\pi(1 - \pi)$ for costs in that spectrum. However, a mechanism splitting the costs has to be chosen for $c \in [\bar{c}_1, \bar{c}(2, 0, \pi)]$, because it is the only mechanism satisfying incentive compatibility and budget constraints.

An ex post efficiency mechanism would provide the good if, and only if, the sum of valuations is weakly larger than provision costs. It would, from the ex ante perspective, provide the good with first best probabilities of $\pi^2 + 2\pi(1 - \pi)$ if $c \in [0, \theta_H]$ and π^2 if $c \in [\theta_H, 2\theta_H]$.

We can see that the voluntary contribution mechanism provides the good with first best probabilities for low values and high values of c. The respective mechanism chosen (single payment for low c and cost-splitting for high c) induces a sufficient level of pivotality and the free-riding issue does not pose a problem. However, a single payment should be used from a first best perspective for medium levels of c, but the level of pivotality is not sufficient to implement an individual payment creating aggregate payments that satisfy the budget constraint. The designer is forced to choose the cost-splitting mechanism.

Inefficiencies in the provision probability arise for medium values of costs ($c \in [\bar{c}_1, \theta_H]$), because the single payment mechanism does not induce incentive compatible payments that satisfy the ex ante budget constraint. The first best measure provides the good with probability $\pi^2 + 2\pi(1 - \pi)$ in this range of costs; the voluntary mechanism only with π^2, yielding a difference of $2\pi(1 - \pi)$. We can see that inefficiencies are increasing in uncertainty, i.e. are largest when π is close to 0.5. This is because uncertainty faced by individuals decreases the probability of being pivotal. Uncertainty

increases the free-riding problem in the voluntary contribution mechanism.

A second way of analysing these inefficiencies is by considering the basic first best measure outlined earlier, i.e. provision if $E\left[\sum_{i=1}^{n} \theta_i\right] \geq c$. Such a measure would provide the good as long as $c \in [0, 2\pi\theta]$. One could construe a similar ex ante measure for a voluntary contribution mechanism by quantifying an upper level of costs. The good would then be provided if expected aggregate payments cover costs, $c \in [0, EAP(n, l, \pi)]$, or in terms of the budget constraint: $\bar{c} = EAP(n, l, \pi)$. Essentially, the provision probability would be one whenever expected aggregate payments exceed costs. For this to work, agents must think that the designer uses the provision rule outlined earlier, namely $Q = 1$ if $\#\theta_H \geq n - l$, while the designer actually bases his decision on $Q = 1$ if $EAP(n, l, \pi) \geq c$. If agents knew that this was the designer's provision rule, they would not have an incentive to truthfully reveal their types, because the provision decision has already been made.[12] In this measure, efficiency is determined by the upper level of c for which the mechanism can provide the good. This upper level is purely determined by EAP and a designer would, therefore, opt for a mechanism maximising expected aggregate payments. One might find such a setting purely theoretic and unrealistic. However, we know that in the 'normal' setting, presented above, the budget constraint holds if $c \leq \bar{c} = \frac{EAP(n,l,\pi)}{\Pr(\#\theta_H \geq n-l)}$. After rearranging to $EAP(n, l, \pi) \geq \Pr(\#\theta_H \geq n - l) \times c$ we can see that the budget constraint implies that the expected aggregate payment is a measure of the combination of the costs and provision probability. Mechanisms with high EAP will maximise the trade-off between the level of costs for which they exist and provision probability. In any way, any efficiency measure depends positively on expected aggregate payments. While the mechanism maximising expected aggregate payments is not necessarily optimal in terms of provision probability, it will exist and provide the public good with positive probability for high levels of c.

The analysis has highlighted a feature missing in the existing literature, which focuses largely on asymptotic behaviour of voluntary contribution mechanisms, rather than on efficiency levels for given group sizes. It appears that efficiency can have several dimensions. If presented with a specific level of c and asked which mechanism is more efficient, the answer is straightforward: the mechanism with higher provision probability. However, if asked which mechanism is generally speaking the more efficient one, then one has to take into account that costs can take different values. One has to trade off provision probabilities for different levels of costs, making the mechanism maximising expected aggregate payments attractive.

This result is depicted in figure 1, in which $\theta_H = 1$. The blue line represents expected aggregate

---

[12] For this provision rule to work, one could argue that agents themselves are interested in running a balanced budget in expectation. This could reintroduce the incentive to reveal types.

payments for a single payment mechanism and the red line expected aggregate payments for the cost-splitting mechanism. The dashed lines refer to the upper bound of costs for which these mechanisms exist.
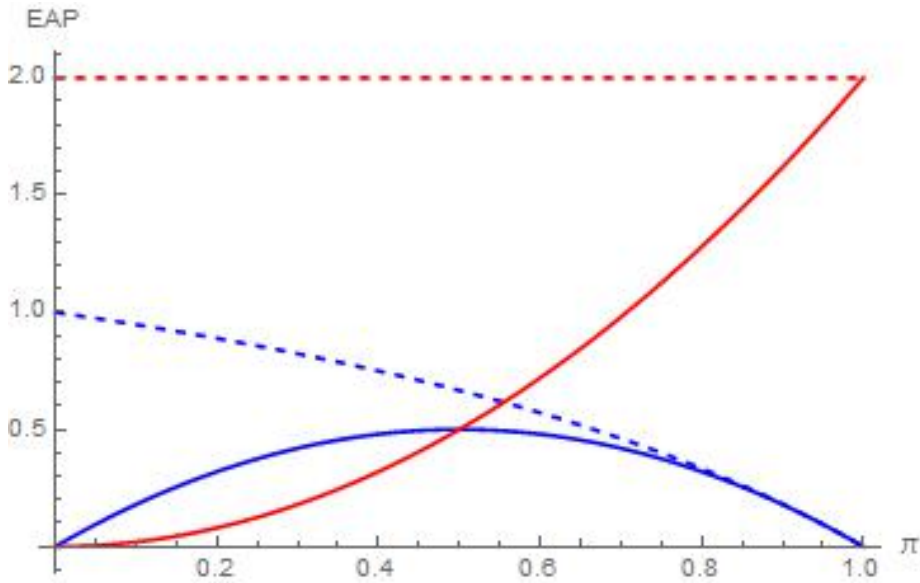

Figure 1: EAP^*, $n = 2$

In terms of provision probability, the designer should use a single payment mechanism as long as the costs lie below the blue dotted line, while a cost-splitting mechanism has to be used if costs lie above it.

2.1.2 N agents

We can expand the logic of the two agent case to derive the constraints, interim payments and expected aggregate payments for the n agent case. Rather than deriving individual and expected aggregate payments as well as provision probabilities separately for each possible mechanism, we can build on the two agent case to derive general incentive compatibility constraints as functions of the chosen mechanism.[13] These can then be used to derive individual and aggregate payments as well as provision probabilities as functions of the chosen mechanism. It will also allow us to find a more specific functional form for the general budget constraint. All of this will then allow us to analyse efficiency levels in the n agents case, to identify mechanisms that maximise expected aggregate payments and to describe the behaviour of voluntary contribution mechanisms for increasing group

---

[13]As shown, the participation constraint is slack, because it is weaker than the incentive compatibility constraint for high types. With our assumption of $\theta_L = 0$ and, therefore, $P_L = 0$, we can solely focus on the incentive compatibility constraint of the high type. It is the only constraint determining interim payments $P_H(n, l, \pi)$.

125

sizes.

The general **incentive compatibility constraint** can be written as:

$$\boldsymbol{\pi_{-i,l}}\Delta \boldsymbol{U_l}(\boldsymbol{\theta_{-i}}) \geq 0$$

with

$$\boldsymbol{\pi^T_{-i,l}} = \begin{pmatrix} \binom{n-1}{n-1}\pi^{n-1}(1-\pi)^0 \\ \binom{n-1}{n-1-1}\pi^{n-1-1}(1-\pi)^1 \\ \vdots \\ \binom{n-1}{n-1-l}\pi^{n-1-l}(1-\pi)^l \\ \binom{n-1}{n-1-(l+1)}\pi^{n-1-(l+1)}(1-\pi)^{l+1} \\ \vdots \\ \binom{n-1}{0}\pi^0(1-\pi)^{n-1} \end{pmatrix}$$

and

$$\Delta \boldsymbol{U_l}(\boldsymbol{\theta_{-i}}) = \begin{pmatrix} U(\theta_H,\boldsymbol{\theta_{n-1}})-U(0,\boldsymbol{\theta_{n-1}}) \\ U(\theta_H,\boldsymbol{\theta_{n-1-1}})-U(0\boldsymbol{\theta_{n-1-1}}) \\ \vdots \\ U(\theta_H,\boldsymbol{\theta_{n-1-l}})-U(0,\boldsymbol{\theta_{n-1-l}}) \\ U(\theta_H,\boldsymbol{\theta_{n-1-(l+1)}})-U(0,\boldsymbol{\theta_{n-1-(l+1)}}) \\ \vdots \\ U(\theta_H,\boldsymbol{\theta_0})-U(0,\boldsymbol{\theta_0}) \end{pmatrix}.$$

Proof. We can write the incentive compatibility constraint for the case of n agents, $EU(Q(\theta_i,\boldsymbol{\theta_{-i}})) \geq EU(Q(\hat{\theta}_i,\boldsymbol{\theta_{-i}}))$, as:

$$\boldsymbol{\pi_{-i}}\boldsymbol{U}(\theta_H,\boldsymbol{\theta_{-i}}) \geq \boldsymbol{\pi_{-i}}\boldsymbol{U}(0,\boldsymbol{\theta_{-i}})$$

where $\boldsymbol{\pi_{-i}}$ is a $1 \times (n-1)$ row-vector of probabilities of how many other agents value the public good. $\boldsymbol{U}(\cdot,\boldsymbol{\theta_{-i}})$ is a $(n-1) \times 1$ column-vector of the utilities depending on how many other agents declare that they value the public good.[14] The agent must find it optimal to reveal his true type θ_H, taking expectations over other agents' valuations.

Subtract the right hand side from the left hand side and use $\Delta \boldsymbol{U}(\boldsymbol{\theta_{-i}})$ for the difference in the agent's utility depending on whether he declares to value the good or not $(\boldsymbol{U}(\theta_H,\boldsymbol{\theta_{-i}}) - \boldsymbol{U}(0,\boldsymbol{\theta_{-i}}))$. The incentive compatibility constraint can then be written as:

$$\boldsymbol{\pi_{-i}}\Delta \boldsymbol{U}(\boldsymbol{\theta_{-i}}) \geq 0.$$

There are n different mechanisms to choose from; a mechanism in which one payment is sufficient to cover the cost, a mechanism in which two payments are needed, up to a mechanism in which n

---

[14] $\boldsymbol{\pi^T_{-i}} = \begin{pmatrix} \binom{n-1}{n-1}\pi^{n-1}(1-\pi)^0 \\ \binom{n-1}{n-2}\pi^{n-2}(1-\pi)^1 \\ \vdots \\ \binom{n-1}{0}\pi^0(1-\pi)^{n-1} \end{pmatrix}$ and $\boldsymbol{U}(\cdot,\boldsymbol{\theta_{-i}}) = \begin{pmatrix} U(\cdot,\boldsymbol{\theta_{n-1}}) \\ U(\cdot,\boldsymbol{\theta_{n-2}}) \\ \vdots \\ U(\cdot,\boldsymbol{\theta_0}) \end{pmatrix}$. The subscript refers to how many other agents declare that they value the public good.

payments are needed. Incentive compatible payments and utilities derived for different realisations of $\boldsymbol{\theta}$ depend on the specified mechanism.

We can say that there will be an interval $(n-1-l) \times P_H(n,l,\pi) < c \leq (n-l) \times P_H(n,l,\pi)$ for any mechanism, where $l \in (0,1,2,...,n-1)$ states how many agents do not value the public good. By choosing an l the designer chooses how to split the costs or which mechanism to use.[15]

Since the incentive compatible payments depend on the way the costs are split (the l chosen), we have to write the vectors of probabilities and utilities as functions of l, yielding the general incentive compatibility constraint above. □

Maximum **incentive compatible individual payments** as a function of the chosen mechanism are:

$$P_H(n,l,\pi) = \binom{n-1}{n-1-l} \pi^{n-1-l}(1-\pi)^l \theta_H.$$

They depend on θ_H and the pivotality probability, i.e. that $n-1-l$ other agents make a payment when $n-l$ are needed. Only then is an individual's agents payment pivotal for provision.[16]

These yield expected aggregate payments of:

$$EAP(n,l,\pi) = n\pi \binom{n-1}{n-1-l} \pi^{n-1-l}(1-\pi)^l \theta_H.$$

Any mechanism l results in a **probability for provision** of the public good of:

$$\Pr(\#\theta_H \geq n-l) = \sum_{i=n-l}^{n} \binom{n}{i} \pi^i (1-\pi)^{n-i}.$$

Using these two results, we can write the **general budget constraint** as:

$$c \leq \bar{c} = \frac{EAP(n,l,\pi)}{\Pr(\#\theta_H \geq n-l)} = \frac{n\pi \binom{n-1}{n-1-l} \pi^{n-1-l}(1-\pi)^l \theta_H}{\sum_{i=n-l}^{n} \binom{n}{i} \pi^i (1-\pi)^{n-i}}.$$

Proof. We can say that:

$$\Delta U(\boldsymbol{\theta_x}) = \begin{cases} -P_H(n,l,\pi) & \text{if } x > n-1-l \\ \theta_H - P_H(n,l,\pi) & \text{if } x = n-1-l \\ -P_H(n,l,\pi) & \text{if } x < n-1-l \end{cases}$$

for any vector $\boldsymbol{\theta_x}$, where x is the number of agents with high valuations.

Given that $n-l$ payments are needed for provision, the agent is only pivotal when $n-1-l$ other agents make payments as well. In the situation in which there are already sufficient payments, the

---

[15] In the two agents case we derived $l^* = 0$ for $c \in [\bar{c}(2,1,\pi), \bar{c}(2,0,\pi)]$ and $l^* = 1$ for $c \in [0, \bar{c}(2,1,\pi)]$ in terms of provision probabilities, while we showed that $EAP(2,0,\pi) > EAP(2,1,\pi)$ for $\pi > 0.5$.

[16] It is easy to verify the payments derived for the two agent case earlier by plugging in $n=2$ and the respective values for l.

good is provided anyways, resulting in $\Delta U = \theta_H - P_H(n, l, \pi) - \theta_H$. The public good is not provided independent of the agent's action in the third scenario, resulting in $\Delta U = -P_H(n, l, \pi) - 0$. The case of $x > n - 1 - l$ occurs with probability $F(n-1) - F(n-1-l)$, where $F(\cdot)$ is the cumulative distribution function. The case of $x < n - 1 - l$ occurs with $F(n-2-l)$ and the case of pivotality with $f(n-1-l)$.

The incentive compatible payment, depending on the mechanism chosen (specified by a level for l) is, therefore, determined by:

$$\Big(F(n-1) - F(n-1-l)\Big)(-P_H(n, l, \pi)) + F(n-2-l)(-P_H(n, l, \pi)) + f(n-1-l)\Big(\theta_H - P_H(n, l, \pi)\Big) \geq 0.$$

After basic simplifications this yields:

$$P_H(n, l, \pi) \leq \frac{f(n-1-l)\theta_H}{F(n-1) - F(n-1-l) + F(n-2-l) + f(n-1-l)}.$$

The probabilities in the denominator add up to one. The maximum incentive compatible interim payment above follows by applying the binomial distribution.

Expected aggregate payments follow directly from the individual payments, because $EAP(n, l, \pi) = n\pi P_H(n, l, \pi)$. At least $n - l$ successes out of n trials, which happen with probability π each, have to occur for the public good to be provided. The provision probability is then a simple application of the binomial distribution. $\qquad\square$

As shown in the two agent case and outlined in the general setup, the provision probability is increasing in l. The designer will want to set a mechanism with an l as high as possible. However, such a mechanism might not satisfy the general budget constraint for a given c.

When asked which mechanism to choose when faced with a specific level of c in order to maximise the provision probability, the designer has to maximise $\Pr(\#\theta_H \geq n - l)$ subject to the budget constraint for given c and n. The solution to this yields the optimal choices of mechanisms and the respective provision probabilities for all levels of c.[17] As in the two agent case, inefficiencies arise because mechanisms yielding a provision probability equal to the one derived under the first best measure do not implement payments large enough to satisfy the budget constraint. The designer is then forced to choose a mechanism with lower l. This problem arises because expected aggregate payments are not large enough for the mechanism with higher l. Incentive compatible individual payments are insufficiently high due to free rider issues.

As pointed out earlier, all efficiency measures depend on aggregate payments. Furthermore, I have shown that the answer to the question which mechanism is preferable or more efficient in general

---

[17]We have derived the results for this explicitly in the two agent case.

involves a trade-off between provision probabilities at all possible levels of c. The mechanism maximising expected aggregate payments seems like a natural candidate to solve this trade-off. We are going to see that this mechanism provides the good for the highest possible $c = \bar{c}$ in expectation in sufficiently large groups. It also provides a further insight, because we know that there exist mechanisms for each $c < \bar{c}$ that will provide the good with a probability larger than 0.5 if the budget constraint is slack under this mechanism. This must be the case, because the designer can choose a mechanism with higher l, resulting in a higher provision probability, if the budget constraint is slack. Finally, maximising EAP is useful for our comparison to the heterogeneous group.

Proposition 1.

*In a homogeneous group of size n, the **mechanism maximising expected aggregate payments**, defined by l^*, is given by the integer value satisfying:*

$$n - n\pi - 1 \leq l^* \leq n - n\pi.$$

*For any combination of n and π, it results in **maximum individual interim payments** of:*

$$P_H^*(n, \pi) = \binom{n-1}{\lfloor n\pi \rfloor} \pi^{\lfloor n\pi \rfloor} (1 - \pi)^{n - \lfloor n\pi \rfloor - 1} \theta_H.$$

Maximum expected aggregate payments *implemented by the mechanism are:*

$$EAP^*(n, \pi) = n\pi P_H^*(n, \pi) = n\pi \binom{n-1}{\lfloor n\pi \rfloor} \pi^{\lfloor n\pi \rfloor} (1 - \pi)^{n - \lfloor n\pi \rfloor - 1} \theta_H.$$

Proof. Since the expected aggregate payments for a specific mechanism determined by l are given by $EAP(n, l, \pi) = n\pi P_H(n, l, \pi)$ and l is the only parameter of choice, maximising individual payments will also result in maximising expected aggregate payments.

The mechanism maximising individual payments and, therefore, also expected aggregate payments is given by:

$$\arg\max_l P_H(n, l, \pi).$$

We are, therefore, looking for the mode of the distribution of the number of other agents that value the good. The mode m of the binomial distribution with $n - 1$ draws and probability π is given by:

$$m(B(n - 1, \pi)) = \begin{cases} \lfloor n\pi \rfloor & \text{if } n\pi \notin \mathbb{Z} \\ n\pi \vee n\pi - 1 & \text{if } n\pi \in \mathbb{Z} \end{cases},$$

which yields the inequality defining the mechanism l^* above.

Maximum individual and aggregate payments are then derived by simply plugging l^* into the general functions derived earlier. Since $\lfloor n\pi \rfloor = n\pi$ when $n\pi$ is an integer, the floor of $n\pi$ is always

a mode. We can work with this and neglect the possibility of the existence of another mode, because that mechanism would yield the same payments. □

Figure 2 highlights the behaviour of the optimal choice of mechanism, if the aim was to maximise EAP, expected aggregate payments themselves and the upper bound on costs for ex ante first best provision. It aims to provide an idea of how aggregate payments behave as the group size increases and how they depend on the choice of the cost-splitting rule.
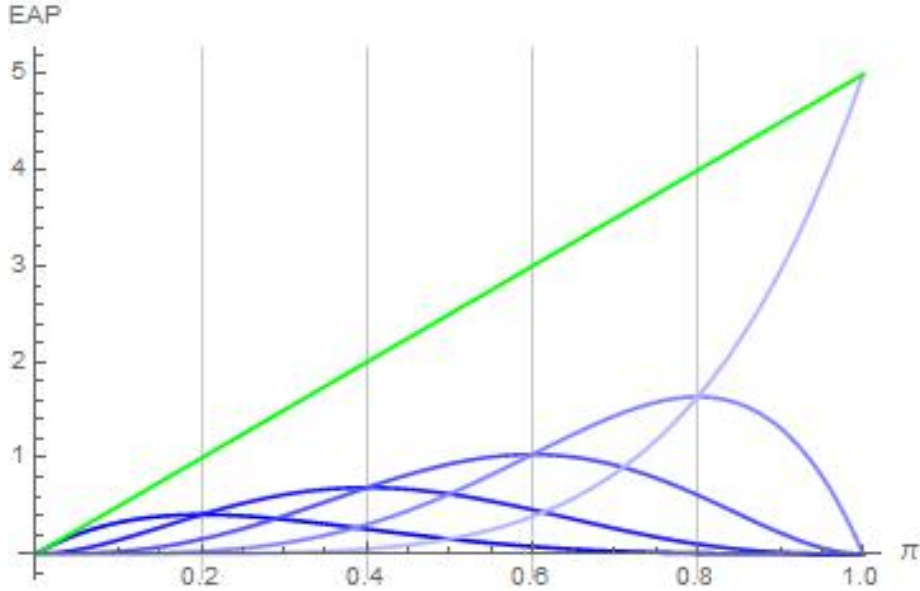

Figure 2: Choice of mechanism, EAP and $n\pi\theta_H$ for $n = 5$

Figure 2 is restricted to $n = 5$ for clarity. The results shown generalise for all n. The blue lines show expected aggregate payments for different choices of mechanism. The dark blue is a mechanism in which a single payment is sufficient to provide the good. The lighter the blue line, the more payments are needed; up to the lightest blue line, which shows the EAP for a mechanism in which five payments are needed for provision. $EAP(n, \pi)^*$, as derived earlier, is just given by the upper envelope of all the expected aggregate payment functions.

Starting from the origin, the optimal mechanism in terms of maximising EAP switches to one where one more payment is needed every time a vertical line is passed. Finally, the green line depicts $5\pi\theta$, our measure of ex ante efficiency: provision if $E[\sum_{i=1}^{n} \theta_i] = n\pi\theta \geq c$. The graph gives a first sight at how expected aggregate payments depend on the choice of mechanism and highlights that $EAP(n, \pi)^*$ is increasing in π if one always chooses the mechanism identified by l^* earlier.

We can use the derived functions to establish further results for large groups in line with the literature in a simpler, more intuitive way. We can also derive additional results.

Proposition 2.

- *individual maximum payments are decreasing in group size (in line with Mailath and Postlewaite (1990), and Hellwig (2003));*

- *maximum expected aggregate payments are increasing in group size (in line with Hellwig (2003));*

- *inefficiencies are increasing in group size (in line with Mailath and Postlewaite, and partially Hellwig (2003)); and*

- *the mechanism maximising expected aggregate payments, defined by l*, will provide the public good in expectation for sufficiently large groups.*

The first and third results are equivalent to Mailath and Postlewaite (1990) and Hellwig (2003), while the second matches Hellwig's (2003) results. The final result further highlights that the voluntary provision of public goods in large groups does not necessarily collapse due to an increased free rider problem. It reinforces Hellwig's (2003) argument.

Proof. I will make use of a normal approximation to the binomial distribution based on Prohorov (1953).[18] The approximations are valid, or reasonably precise, for the usual rule of thumb for normal approximations of binomial distributions, namely roughly $n\pi(1-\pi) > 9$. For simplification purposes I also assume that $n\pi$ and $(n+m)\pi$ are integers.[19]

$$Pr\left[X = x\right] = \frac{1}{\sqrt{n\pi(1-\pi)}}\frac{1}{\sqrt{2\pi}}\exp\left[-\frac{1}{2}\frac{(x-n\pi)}{n\pi(1-\pi)}\right]$$

Evaluating it at the mode, we can see that:

$$\frac{Pr_{n+m}\left[X = \lfloor(n+m)\pi\rfloor\right]}{Pr_n\left[X = \lfloor n\pi\rfloor\right]} = \frac{\sqrt{n}}{\sqrt{n+m}} = \frac{1}{\sqrt{1+\frac{m}{n}}} < 1.$$

This tells us that individual payments are decreasing in group size as the variance and, therefore, the modal probability decreases with a larger group size.

As for the total expected payments, we can say that:

$$\frac{EAP_{n+m}}{EAP_n} = \frac{(n+m)\pi Pr_{n+m}\left[X = \lfloor(n+m)\pi\rfloor\right]}{n\pi Pr_n\left[X = \lfloor n\pi\rfloor\right]} = \sqrt{1+\frac{m}{n}} > 1.$$

---

[18]Any normal approximation would work here. This one was purely chosen for simplicity and precision. In addition, we would usually need a continuity correction when approximating our discrete distribution with a continuous one. The approximation used here approximates the normal approximation with continuity correction.

[19]The result can also be derived for non-integer values of $n\pi$ using an approximation of the ratios of Gamma functions, such as the one in LaForgia and Natalini (2011).

Expected aggregate payments are, therefore, increasing in group size.

Our ex ante efficiency rule behaves as follows:

$$\frac{FB_{n+m}}{FB_n} = \frac{(n+m)\pi\theta}{n\pi\theta} = 1 + \frac{m}{n}.$$

Convergence rates of individual and aggregate payments and the speed of divergence between the voluntary and the first best allocation are of the same order as described in the literature by Hellwig (2003).

The mechanism maximising expected aggregate payments, which as shown is determined by l^*, requires there to be at least $\lfloor n\pi \rfloor$ or $n\pi - 1$ agents reporting high valuations. The provision probability is, hence, given by $\Pr(\#\theta_H \geq \lfloor n\pi \rfloor)$, which is increasing in n and converges to $0.5 + \epsilon$. We have seen that l^* is essentially determined by the expected value. The provision probability is, thus, given by having a number of successes equal to the mean minus a bit/one. Under our second ex ante provision rule of $Q = 1$ if $EAP \geq c$, this mechanism would be most efficient as it maximises the level of c under which the good is provided. This mechanism trading off provision probabilities over different levels for c will provide the good up to $\bar{c}$ in expectation. We can, therefore, also be assured that there exist mechanisms with higher l that provide the good with higher probabilities for all $c < \bar{c}$ when the budget constraint is slack, . $\qquad\square$

Properly proving this behaviour for very small population sizes is not really feasible in this discrete model, although one could establish further quantitative results using a suitable approximation for the modal probability of the binomial distribution. The intuitive explanation for the results is that the variance of the distribution is increasing in the population size and that the modal probability must, therefore, be decreasing. This causes individual payments to decrease. However, the decrease in individual payments is outweighed by the increasing population size, causing expected aggregate payments to increase.

Figures 3 and 4 show these results numerically. $P_H(n,\pi)^*$ are at their minimum for $\pi = 0.5$, because the variance of the binomial distribution is maximised and the modal probability minimised. As seen earlier in the two agent case, this causes the free-riding problem and, hence, inefficiencies to be maximised. The payments are then monotonically increasing in a symmetric way as π either increases or decreases. Finally, individual payments are monotonically decreasing in the group size n. It is straightforward to show analytically that $\frac{\partial P_H(n,\pi)^*}{\partial n} < 0$, $\left.\frac{\partial P_H(n,\pi)^*}{\partial \pi}\right|_{\pi<0.5} < 0$, $\left.\frac{\partial P_H(n,\pi)^*}{\partial \pi}\right|_{\pi=0.5} = 0$, $\left.\frac{\partial P_H(n,\pi)^*}{\partial \pi}\right|_{\pi>0.5} > 0$ and $P_H(n,\pi)^*_{\pi+\epsilon} = P_H(n,\pi)^*_{\pi-\epsilon}$. Furthermore, expected aggregate payments are monotonically increasing in the population size and π; $\frac{\partial EAP(n,\pi)^*}{\partial n} > 0$ and

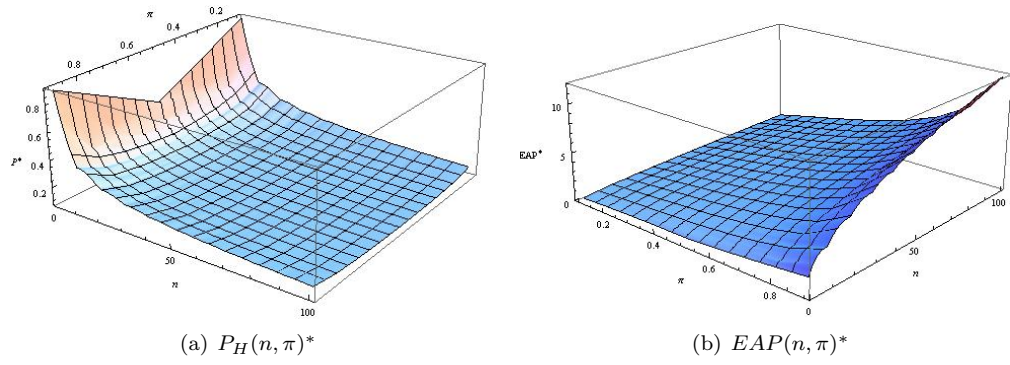$$\frac{\partial EAP(n,\pi)^*}{\partial \pi} > 0.$$

(a) $P_H(n,\pi)^*$

(b) $EAP(n,\pi)^*$

Figure 3: $n \in (2, 100)$, $\pi \in [0, 1]$
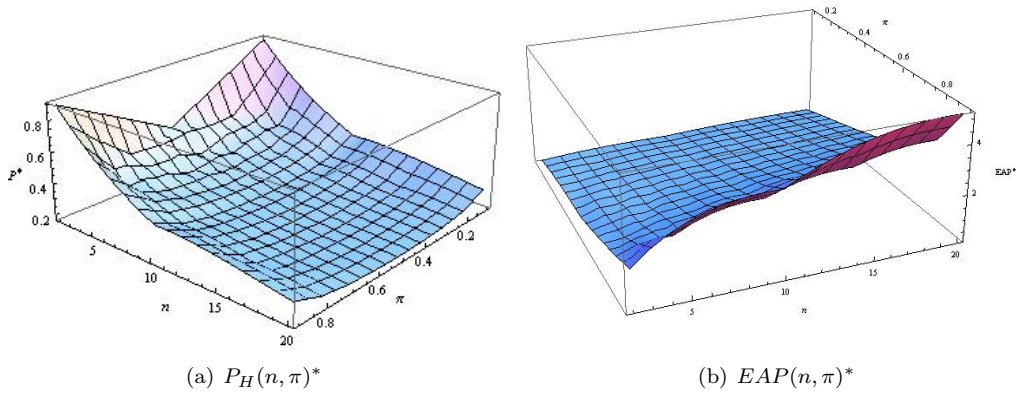

(a) $P_H(n,\pi)^*$

(b) $EAP(n,\pi)^*$

Figure 4: $n \in (2, 20)$, $\pi \in [0, 1]$

Figure 5(a) shows how much higher individual payments are as one moves away from $\pi = 0.5$, which is depicted by the black line. The coloured lines show the payments as we decrease or increase π by 10 percentage points, so that the red line shows the maximum individual payments for $\pi = 0.9$ or $\pi = 0.1$.

Figure 5(b) shows the difference of expected aggregate payments from $\pi = 0.1$ (the black line) up to $\pi = 0.9$ (the red line), while increasing π in steps of 20 percentage points.
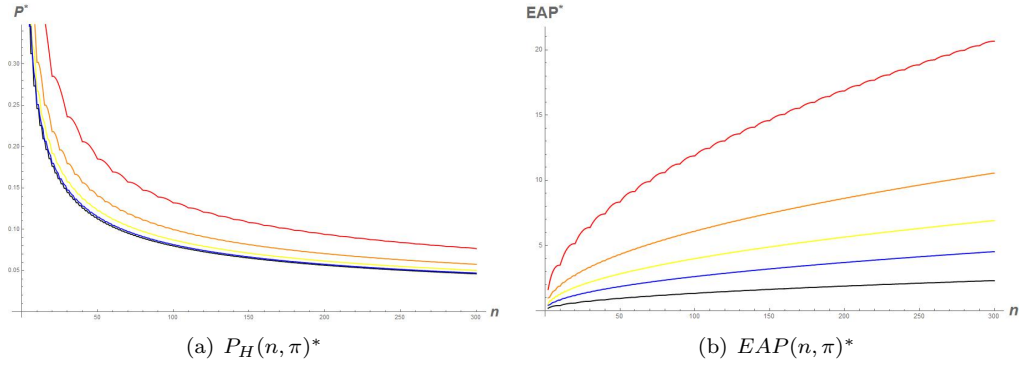
| (a) $P_H(n,\pi)^*$ | (b) $EAP(n,\pi)^*$ |

Figure 5: $n \in (2, 300)$

2.2 Heterogeneous agents

Assume that it is known and observable whether an agent is a member of subpopulation A or B. The designer can, therefore, charge a different price to identical types, as long as they are members of different subpopulations. Members of different subgroups are pooled with a different set of agents. This means that they have non-identical incentive compatibility constraints. This chapter will describe how price discrimination might affect individual and aggregate payments, but will first analyse how payments are different compared to the homogeneous case even when price discrimination is not used. The difference of payments, compared to the homogeneous case, is going to be driven by:

- the distributional effect;
- the group composition effect; and
- the modal effect.

These effects and how they affect individual and aggregate payments will be explained in detail.

2.2.1 No price discrimination

Assume that the designer does not charge different prices and does not use price discrimination between subgroups. The designer is either not able to use price discrimination, or he is, but it is considered unfair to charge agents with identical valuations different prices based on what part of the population they belong to.

In this case, the designer can only specify one payment. Agents in different subgroups, however, face different incentive compatibility constraints. Since there is again just one P_H, payments of members of subgroup A and B are perfect substitutes and identical. The payment will still depend

on a single l specified. The only thing that determines pivotality is the number of payments (the number of agents declaring valuation for the public good).

2.2.1.1 Two Agents

I am going to start with the simple case of two agents in this heterogeneous case, with $n_A = n_B = 1$. We are going to analyse this case in depth and step-by-step to highlight differences to the homogeneous case and improve intuition for the general n agent case. There will be a single payment that is independent of subgroup membership without the use of price discrimination. Therefore, there are again two possible mechanisms. The intuition in both cases is identical to the homogeneous case.

Let us consider the *single payment mechanism*, defined by $l = 1$. Evaluating the incentive compatibility constraints results in the following maximum individual payments for high types:

$$P_{H,A}(\cdot) = (1 - \pi_B)\theta_H$$

$$P_{H,B}(\cdot) = (1 - \pi_A)\theta_H$$

where the subscript refers to subgroup membership.

The designer is limited by the smaller payment when price discrimination is not used. One of the incentive compatibility constraints is slack; in this case it is the incentive compatibility constraint of agent A, because $(1 - \pi_A) < (1 - \pi_B)$. The incentive compatible individual payment is, therefore, given by:

$$P_H(\cdot) = (1 - \pi_A)\theta_H.$$

The single payment mechanism produces expected aggregate payments, $EAP(n_A, n_B, \pi_A, \pi_B, l)$, of:

$$EAP(1, 1, \pi_A, \pi_B, 1) = (\pi_A + \pi_B)(1 - \pi_A)\theta_H$$

a provision probability of:

$$\Pr(\#\theta_H \geq 1) = \pi_A \pi_B + \pi_A(1 - \pi_B) + \pi_B(1 - \pi_A)$$

and satisfies the ex ante budget constraint for:

$$c \leq \bar{c}(1, 1, \pi_A, \pi_B, 1) = \frac{(\pi_A + \pi_B)(1 - \pi_A)}{\pi_A \pi_B + \pi_A(1 - \pi_B) + \pi_B(1 - \pi_A)}\theta_H.$$

However, the initial condition $c \leq P_H$ has to be satisfied for the single payment mechanism to exist. It is straightforward to show that this constraint is stricter than the budget constraint

$(1 - \pi_A < \bar{c}(1, 1, \pi_A, \pi_B, 1))$. The single payment mechanism does, therefore, only exist for:

$$c \le (1 - \pi_A)\theta_H = \bar{c}_{1,h}.$$

Let us now consider the *cost-splitting mechanism*, defined by $l = 0$. Evaluating the incentive compatibility constraints for the cost-splitting mechanism gives us the following maximum individual payments:

$$P_{H,A}(\cdot) = \pi_B \theta_H$$

$$P_{H,B}(\cdot) = \pi_A \theta_H.$$

We are again limited by the smaller payment and one of the incentive compatibility constraints is slack, because $\pi_A < \pi_B$ and the designer does not make use of price discrimination. In this case, it is the constraint of agent B, because $\pi_A > \pi_B$. The incentive compatible payment without price discrimination for the cost-splitting mechanism is, therefore, given by:

$$P_H(\cdot) = \pi_B \theta_H.$$

This mechanism produces expected aggregate payments of:

$$EAP(1, 1, \pi_A, \pi_B, 0) = (\pi_A + \pi_B)\pi_B \theta_H$$

a provision probability of:

$$\Pr(\#\theta_H \ge 2) = \pi_A \pi_B$$

and satisfies the ex ante budget constraint for:

$$c \le \bar{c}(1, 1, \pi_A, \pi_B, 0) = \left(1 + \frac{\pi_B}{\pi_A}\right)\theta_H.$$

The basic principle that it is optimal to set a mechanism with l as high as feasible, where feasibility is given by the ex ante budget constraint, still holds. The optimal mechanism in order to maximise provision probability is $l^* = 1$ with its associated payments for $c \in [0, \bar{c}_{1,h}]$ and $l^* = 0$ with its associated payments for $c \in [\bar{c}_{1,h}, \bar{c}(1, 1, \pi_A, \pi_B, 0)]$.

Proposition 3.

Comparing a two agent group of heterogeneous agents, with $n_A = n_B = 1$ and probabilities of π_A and π_B, to the equivalent homogeneous group, with $n = 2$ and $\pi = \frac{\pi_A + \pi_B}{2}$, inefficiencies are larger in the heterogeneous groups when no price discrimination is used.

Proof. Using the same first best measure as in the homogeneous case, we can say that the good is provided with first best probabilities of $\pi_A \pi_B + \pi_A(1 - \pi_B) + \pi_B(1 - \pi_A)$ if $c \in [0, \theta_H]$ and

$\pi_A\pi_B$ if $c \in [\theta_H, 2\theta_H]$. As in the homogeneous case, inefficiencies arise for medium cost levels ($c \in [\bar{c}_{1,h}, \theta_H]$), because the single payment mechanism does not create incentive compatible payments large enough to satisfy the ex ante budget constraint. Additional inefficiencies arise for $c \in [\bar{c}(1, 1, \pi_A, \pi_B, 0), 2\theta_H]$, because there exists no mechanism satisfying the ex ante budget constraint for high costs when price discrimination is not used.

We can compare efficiency in the homogeneous and heterogeneous case in several ways. First, we have seen that inefficiencies arise in the homogeneous case for $c \in [(1 - \pi)\theta_H, \theta_H]$, while they arise for $c \in [(1 - \pi_A)\theta_H, \theta_H]$ in the heterogeneous case. Comparing equivalent cases by using $\pi = \frac{\pi_A + \pi_B}{2}$ in the homogeneous case, and ensuring that both cases yield the same first best provision rule, we can see that the set of costs for which the single payment does not satisfy the budget constraint is larger in the heterogeneous case. The designer has to use the less efficient cost-splitting mechanism more often in the heterogeneous setting without price discrimination. Adding this to the fact that there exists an entire other set of costs in the heterogeneous setting for which no mechanism satisfies the budget constraint (at high c), shows that the set of costs for which an inefficient mechanism has to be used is larger in the heterogeneous setting with two agents.

Not only is the set of costs for which there are inefficiencies larger, but the inefficiencies measured as the difference in provision probabilities are also larger in the heterogeneous setting. This is verified easily by using $\pi = \frac{\pi_A + \pi_B}{2}$ and showing that $\pi_A(1 - \pi_B) + \pi_B(1 - \pi_A) > (\pi_A + \pi_B)\left(1 - \frac{\pi_A + \pi_B}{2}\right)$ for $\pi_A > \pi_B$ and $1 \geq \pi_j \geq 0$. The proposition follows immediately from this when inefficiencies are measured by provision probabilities.

Using the second way of assessing efficiency, in which we measure efficiency purely by the level of expected aggregate payments, it is easy to show that maximised expected aggregate payments are always smaller in the heterogeneous case when compared to the equivalent homogeneous case. This means that inefficiencies are larger in the heterogeneous case, independent of how we measure inefficiencies. $\qquad\square$

Intuitively these inefficiency results for the small heterogeneous case stem from the fact that price discrimination is not used, rather than from distributional effects. The lack of price discrimination means that only one agent's incentive compatibility constraint is binding. The designer does not extract the maximum payments possible. It is not clear whether these inefficiency results for the heterogeneous case without price discrimination hold for large groups as well, or whether statistical/distributional effects will affect individual and aggregate payments and potentially dominate this effect in larger groups.

2.2.1.2 N Agents

The maximum incentive compatible payment with heterogeneous agents and no price discrimination is given by:

$$P_H(\cdot, l) = \min[f_A(n-1-l)\theta_H, f_B(n-1-l)\theta_H].$$

The provision probability is, once again, given by:

$$\Pr(\#\theta_H \geq n - l)$$

which is, however, not just determined by the simple cumulative distribution of a Binomial distribution as in the homogeneous case. Instead, the distribution is a convolution of n_A Bernoulli trials with probability π_A and n_B Bernoulli trials with probability π_B.

Finally, the mechanism specified by l yields expected aggregate payments of:

$$EAP(\cdot, l) = (n_A\pi_A + n_B\pi_B) \times \min[f_A(n-1-l)\theta, f_B(n-1-l)\theta]^{20}$$

and satisfies the ex ante budget constraint if:

$$c \leq \bar{c} = \frac{EAP(\cdot, l)}{\Pr(\#\theta_H \geq n - l)}.$$

Proof. The column-vector $\Delta U(\boldsymbol{\theta_{-i}})$ is unchanged from the homogeneous case, because payments are perfect substitutes. However, the row-vector of probabilities $\left(\boldsymbol{\pi_{-i}}\right)$ is changed and depends on which subpopulation the agent belongs to:

$$\boldsymbol{\pi_{A,-i}} = \begin{pmatrix} \binom{n-1}{n-1}\pi_A^{n_A-1}(1-\pi_A)^0\pi_B^{n_B}(1-\pi_B)^0 \\ \sum_{i=0}^{1}\binom{n_A-1}{n_A-1-i}\pi_A^{n_A-1-i}(1-\pi_A)^i\binom{n_B}{n_B-1+i}\pi_B^{n_B-1+i}(1-\pi_B)^{1-i} \\ \sum_{i=0}^{2}\binom{n_A-1}{n_A-1-i}\pi_A^{n_A-1-i}(1-\pi_A)^i\binom{n_B}{n_B-2+i}\pi_B^{n_B-2+i}(1-\pi_B)^{2-i} \\ \vdots \\ \binom{n-1}{0}\pi_A^0(1-\pi_A)^{n_A-1}\pi_B^0(1-\pi_B)^{n_B} \end{pmatrix}$$

and

$$\boldsymbol{\pi_{B,-i}} = \begin{pmatrix} \binom{n-1}{n-1}\pi_A^{n_A}(1-\pi_A)^0\pi_B^{n_B-1}(1-\pi_B)^0 \\ \sum_{i=0}^{1}\binom{n_A}{n_A-i}\pi_A^{n_A-i}(1-\pi_A)^i\binom{n_B-1}{n_B-2+i}\pi_B^{n_B-2+i}(1-\pi_B)^{1-i} \\ \sum_{i=0}^{2}\binom{n_A}{n_A-i}\pi_A^{n_A-i}(1-\pi_A)^i\binom{n_B-1}{n_B-3+i}\pi_B^{n_B-3+i}(1-\pi_B)^{2-i} \\ \vdots \\ \binom{n-1}{0}\pi_A^0(1-\pi_A)^{n_A}\pi_B^0(1-\pi_B)^{n_B-1} \end{pmatrix}.$$

The probabilities of being pivotal depend on the specific mechanism, which, without price discrimination, is still just identified by a single l, because payments from members of different subgroups

---

[20]For a homogeneous and heterogeneous group to be comparable it has to be that $\pi = \bar{\pi} = \frac{n_A\pi_A + n_B\pi_B}{n_A + n_B}$, so that the two yield the same first best results. For $\pi = \bar{\pi}$, we can see that $n\bar{\pi} = n_A\pi_A + n_B\pi_B$. It is, thus, sufficient to check in which case individual payments are larger in order to compare expected aggregate payments for equivalent groups.

are perfect substitutes. We can, therefore, rewrite the pivotality probabilities as functions of the specified cost-splitting rule:

$$
\boldsymbol{\pi}_{A,l,-i} = \begin{pmatrix}
\binom{n-1}{n-1}\pi_A^{n_A-1}(1-\pi_A)^0\pi_B^{n_B}(1-\pi_B)^0 \\
\sum_{i=0}^{1}\binom{n_A-1}{n_A-1-i}\pi_A^{n_A-1-i}(1-\pi_A)^i\binom{n_B}{n_B-1+i}\pi_B^{n_B-1+i}(1-\pi_B)^{1-i} \\
\vdots \\
\sum_{i=0}^{l}\binom{n_A-1}{n_A-1-i}\pi_A^{n_A-1-i}(1-\pi_A)^i\binom{n_B}{n_B-l+i}\pi_B^{n_B-l+i}(1-\pi_B)^{l-i} \\
\sum_{i=0}^{l+1}\binom{n_A-1}{n_A-1-i}\pi_A^{n_A-1-i}(1-\pi_A)^i\binom{n_B}{n_B-(l+1)+i}\pi_B^{n_B-(l+1)+i}(1-\pi_B)^{l+1-i} \\
\vdots \\
\binom{n-1}{0}\pi_A^0(1-\pi_A)^{n_A-1}\pi_B^0(1-\pi_B)^{n_B}
\end{pmatrix}
$$

and

$$
\boldsymbol{\pi}_{B,l,-i} = \begin{pmatrix}
\binom{n-1}{n-1}\pi_A^{n_A}(1-\pi_A)^0\pi_B^{n_B-1}(1-\pi_B)^0 \\
\sum_{i=0}^{1}\binom{n_A}{n_A-i}\pi_A^{n_A-i}(1-\pi_A)^i\binom{n_B-1}{n_B-2+i}\pi_B^{n_B-2+i}(1-\pi_B)^{1-i} \\
\vdots \\
\sum_{i=0}^{l}\binom{n_A}{n_A-i}\pi_A^{n_A-i}(1-\pi_A)^i\binom{n_B-1}{n_B-1-l+i}\pi_B^{n_B-1-l+i}(1-\pi_B)^{l-i} \\
\sum_{i=0}^{l+1}\binom{n_A}{n_A-i}\pi_A^{n_A-i}(1-\pi_A)^i\binom{n_B-1}{n_B-1-(l+1)+i}\pi_B^{n_B-1-(l+1)+i}(1-\pi_B)^{l+1-i} \\
\vdots \\
\binom{n-1}{0}\pi_A^0(1-\pi_A)^{n_A}\pi_B^0(1-\pi_B)^{n_B-1}
\end{pmatrix}.
$$

Using the same approach as in the homogeneous case, we can pin down individual interim payments by establishing the case in which the agent is pivotal and its associated probability and get:

$$
\Delta u(\boldsymbol{\theta_x}) = \begin{cases}
-P_H(\cdot) & \text{if } x > n-1-l \\
\theta - P_H(\cdot) & \text{if } x = n-1-l \\
-P_H(\cdot) & \text{if } x < n-1-l
\end{cases}.
$$

Individual maximum payments given a cost-splitting rule are different, because members of different subgroups face different distribution functions (f_A and f_B respectively). Using the same steps as in the homogeneous case, maximum payments are given by each agent's probability of being pivotal:

$$
P_{H,A} \leq f_A(n-1-l)\theta_H
$$

and

$$
P_{H,B} \leq f_B(n-1-l)\theta_H.
$$

In order to maintain the incentive compatibility constraints without price discrimination, the payment is determined by the less strict inequality. Payments for high types of one subgroup are such that their incentive compatibility constraint is slack. This yields the maximum payment presented above. Provision probabilities, expected aggregate payments and the budget constraint follow directly from individual payments. □

Once again, the provision probability maximising mechanism maximises $\Pr(\#\theta_H \geq n-l)$ with respect to l subject to the budget constraint. For reasons outlined earlier, this mechanism is only optimal for a specific c. Which mechanism can be called most efficient is not, in general, clear-

cut. As shown earlier, a mechanism maximising expected aggregate payments allows for expected provision at the highest possible costs and weighs off provision costs and provision probabilities; it maximises a combined measure of the two. I will, therefore, determine the same mechanism in this heterogeneous setting and compare it to the homogeneous case.

Proposition 4.

The mechanism maximising expected aggregate payments will create larger expected aggregate payments for heterogeneous groups for sufficiently large group sizes, even in the absence of price discrimination.

This result, which suggests that voluntary public good provision is more efficient in heterogeneous groups, can be explained by the existence of three individual effects that explain differences in individual and aggregate payments in the two settings:

1. The **distributional effect**. The distribution of the number of agents valuing the public good is Binomial in the homogeneous case, but Poisson-Binomial in the heterogeneous case. This effect will not exist when subgroups have only one member, but will increase individual and aggregate payments otherwise. The effect is increasing in group size.

2. The **group composition effect**. The average probability for valuing the public good of the rest of the group, which is what an agent faces under his incentive compatibility constraint, differs to the average probability of the overall group in the heterogeneous case. This is not the case in the homogeneous setting. The direction of this effect is not clear, but the size of the effect is decreasing in group size.

3. The **modal effect**. The mode of the heterogeneous group overall might differ from the mode of the equivalent homogeneous group. In the heterogeneous setting, modes might also differ for the distributions that members of different subgroups face in their respective incentive compatibility constraints. This effect decreases individual payments. The magnitude of the effect is decreasing in group size.

Proof. Only a single payment, $P_{H,A} = P_{H,B} = P_H$, is set when price discrimination is not used. From an agent's perspective, the probability of being pivotal depends only on the total number of

agents declaring valuation. The probabilities of being pivotal are given by:

$$f_A(n-1-l) = \sum_{i=0}^{n_A+n_B-1-l} \binom{n_A-1}{i}\binom{n_B}{n_A+n_B-1-l-i} \times$$
$$\pi_A^i(1-\pi_A)^{n_A-1-i}\pi_B^{n_A+n_B-1-l-i}(1-\pi_B)^{n_B-(n_A+n_B-1-l)+i}$$

and

$$f_B(n-1-l) = \sum_{i=0}^{n_A+n_B-1-l} \binom{n_A}{i}\binom{n_B-1}{n_A+n_B-1-l-i} \times$$
$$\pi_A^i(1-\pi_A)^{n_A-i}\pi_B^{n_A+n_B-1-l-i}(1-\pi_B)^{n_B-1-(n_A+n_B-1-l)+i}.$$

The pivotality probabilities are non-identical for the two groups. It is much harder to derive the optimal mechanism, individual payments and aggregate payments, because the total number of agents valuing the good is no longer binomially distributed. The distribution for a member of subgroup A is the convolution of $n_A - 1$ Bernoulli trials with probability π_A and n_B Bernoulli trials with probability π_B. For a member of subgroup B, it is the convolution of n_A Bernoulli trials with probability π_A and $n_B - 1$ Bernoulli trials with probability π_B.

Such a distribution, which is essentially a Binomial with non-constant probabilities, is sometimes referred to as the Poisson-Binomial distribution. As in the homogeneous case, the designer is looking to maximise pivotality when trying to maximise agents' individual payments in order to maximise aggregate payments. We, therefore, need to identify the mode of the distribution of the number of valuations and split costs accordingly.

The Poisson-Binomial distribution has the following properties (Darroch, 1964; Samuels, 1965 and Hoeffding, 1956):

- Probabilities first increase convexly, then concavely, then decrease concavely and then convexly. The distribution of the number of successes can be described as bell shaped.
- The mean is given by $\mu = \sum_{i=1}^{n} \pi_i$.
- The mode of the distribution is determined by the mean and/or one or both of the adjacent integers. To be more precise, the mode, m, is characterised as follows (Darroch, 1964):

$$m(B(n_A, n_B, \pi_A, \pi_B)) = \begin{cases} \mu & \text{if } \mu = \lfloor\mu\rfloor \\ \lfloor\mu\rfloor & \text{if } 0 \leq \mu - \lfloor\mu\rfloor < \frac{1}{\lfloor\mu\rfloor+2} \\ \lfloor\mu\rfloor \vee \wedge \lfloor\mu\rfloor + 1 & \text{if } \frac{1}{\lfloor\mu\rfloor+2} \leq \mu - \lfloor\mu\rfloor \leq \frac{n-\lfloor\mu\rfloor}{n-\lfloor\mu\rfloor+1} \\ \lfloor\mu\rfloor + 1 & \text{if } \frac{n-\lfloor\mu\rfloor}{n-\lfloor\mu\rfloor+1} < \mu - \lfloor\mu\rfloor < 1 \end{cases}.$$

- The variance of the distribution is given by $\sigma^2 = \sum_{i=1}^{n} \pi_i(1-\pi_i)$.

First, let us have a look at how the different distribution affects the individual incentive compatibility constraint and maximum payments. We will refer to this effect as the **distributional effect**.

For now, abstract from possible differences in the mean and mode. Only compare the variation in the distributions. With $\overline{\pi} = \frac{n_A \pi_A + n_B \pi_B}{n_A + n_B}$, the aforementioned variance of the Poisson-Binomial with population size n, consisting of two subgroups with sizes n_A and n_B, can be written as:

$$\sigma^2 = n\overline{\pi}(1 - \overline{\pi}) - n_A(\pi_A - \overline{\pi})^2 - n_B(\pi_B - \overline{\pi})^2.$$

The first term is the variance of an equivalent homogeneous population, i.e. a homogeneous population in which the expected value of total valuation is the same.[21] The second and third term refer to the heterogeneity of the groups.

The difference in variances for a homogeneous population and a population consisting of two subgroups with identical means is:

$$Var(n, \overline{\pi}) - Var(n_A + n_B, \pi_A, \pi_B) = \frac{n_A n_B (\pi_A - \pi_B)^2}{n_A + n_B}.$$

Let us take a look at the behaviour of this difference by using our condition for an equivalent homogeneous group: $\overline{\pi} = \frac{n_A \pi_A + n_B \pi_B}{n_A + n_B}$. Solving it for π_B and plugging it into the equation for the variance we get:

$$\Delta Var = \frac{n_A n_B (\pi_A - \frac{(n_A + n_B)\overline{\pi} - n_A \pi_A}{n_B})^2}{n_A + n_B}.$$

Taking basic derivatives, it is easy to verify that:

$$\frac{\partial \Delta Var}{\partial \pi_A} = \frac{2 n_A n_B (1 + \frac{n_A}{n_B})(\pi_A - \frac{(n_A + n_B)\overline{\pi} - n_A \pi_A}{n_B})}{n_A + n_B} > 0$$

$$\frac{\partial^2 \Delta Var}{\partial \pi_A^2} = \frac{2 n_A n_B (1 + \frac{n_A}{n_B})^2}{n_A + n_B} > 0.$$

Increasing π_A while holding the group sizes and the average probability of valuation constant, or, in other words, increasing the between group heterogeneity, decreases the variation in the distribution of the total number of agents that value the good. The decrease is larger, the larger between group heterogeneity, i.e. variation is a concave function of between group heterogeneity. In other words, the difference in variation is a convex function of between group heterogeneity.

Once again ensuring that we do not change the equivalent homogeneous population by plugging $n_B = \frac{n_A(\pi_A - \overline{\pi})}{\overline{\pi} - \pi_B}$ into the equation for the difference of variances, we obtain:

$$\Delta Var = \frac{n_A(\frac{n_A(\pi_A - \overline{\pi})}{\overline{\pi} - \pi_B})(\pi_A - \pi_B)^2}{n_A + \frac{n_A(\pi_A - \overline{\pi})}{\overline{\pi} - \pi_B}}$$

with

$$\frac{\partial \Delta Var}{\partial n_A} = (\pi_A - \pi_B)(\pi_A - \overline{\pi}) > 0.$$

This shows us that the difference in the variances of the two distributions is increasing in population

size. Hence, assuming that modes and means of the distributions $f_A(n-1-l)$ and $f_B(n-1-l)$ are identical and equal to the mode and mean of $f(n-1-l)$, the distributional effect has shown us that modal probabilities, and therefore payments, will be higher in the heterogeneous case, because the variation of a Poisson-Binomial distribution is smaller than the one of a standard Binomial distribution. This effect is increasing in heterogeneity and group size. Clearly, we do not have the full picture yet, as we had established earlier that the homogeneous group outperformed the heterogeneous group in the two agent case.

First, note that we have compared the distributions with n agents. The relevant distributions determining payments are the ones faced in the incentive compatibility constraints, i.e. the probability of the number of $n-1$ other agents valuing the good. We will consider this later on, but would suspect the effect to be close to the effect on the overall distribution, which we just analysed, at least for larger groups.

To draw final conclusions about the modal probability, which ultimately determines payments, we also have to consider that the means and the mode might be different across groups and different to the equivalent homogeneous group. We have ruled this out in our analysis of the distributional effect.

We have seen that $n\pi$ is either the unique mode or one part of the pair of adjacent integers that are both modes of the Binomial distribution. Furthermore, we know from Darroch's (1964) description of the Poisson-Binomial distribution that the mode is identical to the mean when the mean is an integer. The mode can only be unique and equal to $\lfloor \mu \rfloor$ for a small mean and, therefore, for a small group size, because $\lim_{n\to\infty} \frac{1}{\lfloor \mu \rfloor + 2} = 0$. The same argument is true for a unique mode of $\lfloor \mu \rfloor + 1$, because $\lim_{n\to\infty} \frac{n-\lfloor \mu \rfloor}{n-\lfloor \mu \rfloor + 1} = 1$. For larger populations, the mode is either one of them or both at the same time and there is no analytical solution. Only numerical methods can give us the answer.

To simplify our analysis, I will assume that the mode is always given by $\lfloor \mu \rfloor$. Although that might be incorrect at some combinations of probabilities and group sizes, the derived results still hold as long as the overall pattern is clear. Whether the modes of the different distributions are different at any combination of group sizes and heterogeneity depends on the size of the non-integer part of the mean and on whether the difference in probabilities is sufficient to cause different integer parts of the mean. Hence, a difference in modes occurs more often when heterogeneity is large, but not in a systematic way.[22] Our assumption is, thus, correct more often when heterogeneity is low than when it is high. In any case, it is sometimes correct and it is not false in a systematic way; it will be correct in some cases for small groups, for medium sized groups and for large groups. The

---

[22]The occurrence of different modes is uncorrelated with population size.

behaviour of payments as the group sizes increases is still identified correctly.

Define μ_H as the mean of the equivalent homogeneous group, μ_A as the mean faced by a member of group A in the heterogeneous setting and μ_B as the mean faced by a member of group B in the heterogeneous setting:

$$\mu_H = (n-1)\frac{n_A \pi_A + n_B \pi_B}{n_A + n_B}$$

$$\mu_A = (n_A - 1)\pi_A + n_B \pi_B$$

$$\mu_B = n_A \pi_A + (n_B - 1)\pi_B,$$

When comparing the three means of the distributions that agents face in their according incentive compatibility constraints, we can say that:

$$\mu_B > \mu_H > \mu_A, \text{ for } \pi_A > \pi_B.$$

Again, assume that μ_A is an integer for simplification purposes. As before, this will not be the case for most combinations of group sizes and probabilities, but the same argument given earlier applies. Combining our assumptions, we can say that all incentive compatibility constraints are maximised under the same rule of cost splitting if $\mu_A \in \mathbb{Z}$ and thus $\lfloor \mu_B \rfloor = \lfloor \mu_H \rfloor = \mu_A$. In that case, it is solely the relative size of the variances of the distributions faced in the incentive compatibility constraints, summarised by the distributional effect, that determines when payments are larger.

Having described simplifying assumptions, we can analyse the distributional effect in a slightly different way, allowing for the fact that we should not be comparing group sizes and distributions using n and $n_A + n_B$, but group sizes and distributions agents face under their incentive compatibility constraints. This will give us more precise results for the distributional effect.

The variance of the distribution f_A, i.e. the distribution faced by a member of subgroup A in the heterogeneous setting, is given by:

$$Var_{-A} = (n-1)\overline{\pi}_{-A}(1 - \overline{\pi}_{-A}) - (n_A - 1)(\pi_A - \overline{\pi}_{-A})^2 - n_B(\pi_B - \overline{\pi}_{-A})^2$$

where the subscript $-A$ indicates that a single member of subgroup A has been removed.[23]

Equivalently, the variance of the distribution f_B is given by:

$$Var_{-B} = (n-1)\overline{\pi}_{-B}(1 - \overline{\pi}_{-B}) - n_A(\pi_A - \overline{\pi}_{-B})^2 - (n_B - 1)(\pi_B - \overline{\pi}_{-B})^2.$$

Finally, the variance of the distribution faced by a member i of an equivalent homogeneous group

---

[23]Remember that a member of subgroup A faces only $n_A - 1$ members of subgroup A but n_B members of subgroup B.

is:

$$Var_{-i} = (n-1)\frac{n_A\pi_A + n_B\pi_B}{n_A + n_B}\left(1 - \frac{n_A\pi_A + n_B\pi_B}{n_A + n_B}\right).$$

Using basic simplification, the differences in variances ($\Delta Var_{-A} = Var_{-i} - Var_{-A}$ and $\Delta Var_{-B} = Var_{-i} - Var_{-B}$) are as follows:

$$\Delta Var_{-A} = (n-1)\left(\overline{\pi}(1-\overline{\pi}) - \overline{\pi}_{-A}(1-\overline{\pi}_{-A})\right) + \frac{(n_A-1)n_B(\pi_A - \pi_B)^2}{n_A - 1 + n_B}$$

and

$$\Delta Var_{-B} = (n-1)\left(\overline{\pi}(1-\overline{\pi}) - \overline{\pi}_{-B}(1-\overline{\pi}_{-B})\right) + \frac{n_A(n_B-1)(\pi_A - \pi_B)^2}{n_A - 1 + n_B}.$$

The second part of the sum on the right hand side is the **more accurate distributional effect**, using the actual distributions different agents face. Using basic derivatives, we find that:

$$\frac{\partial \Delta Var_{-A}}{\partial \pi_A} = \frac{2(n_A-1)n_B(1 + \frac{n_A}{n_B})(\pi_A - \pi_B)}{n_A + n_B - 1} \geq 0$$

$$\frac{\partial^2 \Delta Var_{-A}}{\partial \pi_A^2} = \frac{2(n_A-1)n_B(1 + \frac{n_A}{n_B})^2}{n_A + n_B - 1} \geq 0$$

$$\frac{\partial \Delta Var_{-A}}{\partial n_A} = \frac{(\pi_A - \pi_B)^2(\pi_A - \overline{\pi})(n_A^2(\pi_A - \pi_B) - \pi_B + 2n_A(\pi_B - \overline{\pi}) + \overline{\pi})}{(n_A(\pi_A - \pi_B) + \pi_B - \overline{\pi})^2} \geq 0$$

$$\frac{\partial \Delta Var_{-B}}{\partial \pi_B} = \frac{2n_A(n_B-1)(-1 - \frac{n_A}{n_B})(\pi_A - \pi_B)}{n_A + n_B - 1} \leq 0$$

$$\frac{\partial^2 \Delta Var_{-B}}{\partial \pi_B^2} = \frac{2n_A(n_B-1)(-1 - \frac{n_A}{n_B})^2}{n_A + n_B - 1} \geq 0$$

$$\frac{\partial \Delta Var_{-B}}{\partial n_B} = -\frac{(\pi_A - \pi_B)^2(\pi_B - \overline{\pi})((n_B-1)^2\pi_A - n_B^2\pi_B + 2n_B\overline{\pi} - \overline{\pi})}{((n_B-1)\pi_A - n_B\pi_B + \overline{\pi})^2} \geq 0$$

We have shown that the distributional effect decreases variation in the heterogeneous case (increases the difference in the variation). It increases the modal probability and maximum individual payments for every agent i, unless i is the sole member of a group. This partially explains our inefficiency result in the two agent case. The positive effect of increased between group heterogeneity on individual payments is convex. According to the distributional effect, payments in the heterogeneous case become relatively larger compared to the equivalent homogeneous case as the population size increases.

Consider now the first term affecting the differences of variances. I will call this the **group composition effect**. We know that:

$$\overline{\pi} > \overline{\pi}_{-A}$$

and

$$\overline{\pi} < \overline{\pi}_{-B},$$

because $\pi_A > \overline{\pi}$ and $\pi_B < \overline{\pi}$. Define ϵ as the decrease in average probability caused by the removal of one agent of subgroup A. Clearly, ϵ is decreasing in group size and increasing in the difference of probabilities $(\pi_A - \pi_B)$. For the incentive compatibility constraints of a member of subgroup A and a member of the equivalent homogeneous group to be equivalent, it must be that $\overline{\pi}_{-A} - \epsilon = \overline{\pi}$. Define η as the increase caused by the removal of a member of subgroup B, with $\overline{\pi}_{-B} + \eta = \overline{\pi}$ for equivalence.[24]

We can use this to write the group composition effect, the first term in the difference of the variances, as:

$$\Delta Var_{-A,GC} = (n-1)\epsilon \left[1 + \epsilon - 2\overline{\pi}\right]$$

and

$$\Delta Var_{-B,GC} = (n-1)\eta \left[2\overline{\pi} - 1 + \eta\right].$$

Since $\eta = \frac{n_A}{n_B}\epsilon$, we can rewrite the group composition effect for a member of group B as:

$$\Delta Var_{-B,GC} = (n-1)\frac{n_A}{n_B}\epsilon \left[2\overline{\pi} - 1 + \frac{n_A}{n_B}\epsilon\right].$$

Only the sign of the term within square brackets is unclear. The overall sign of the group composition effect depends on the parameters within the square brackets. Whenever the average probability in the total population is sufficiently small, $\overline{\pi} \leq \frac{1+\epsilon}{2}$, the group composition effect has the same sign for members of group A as the distributional effect and increases payments compared to the equivalent homogeneous population. The same is true for members of group B whenever the average probability in the total population is sufficiently large, $\overline{\pi} \geq \frac{1-\frac{n_A}{n_B}\epsilon}{2}$.

The intuition is that the removal of agent i in the homogeneous case does not affect the average probability in the rest of the population. The variance of the distribution in the incentive compatibility constraint is smaller than in the whole population simply because there is one less Bernoulli draw. In the heterogeneous case, the average probability for the rest of the population is different and differs for members of different subgroups. The variance is, therefore, further decreasing compared to the homogeneous case when the average probability of the distribution faced under the incentive compatibility constraint is further away from 0.5 than the average probability in the whole population. This means that the group composition effect increases payments for all agents as long as $\frac{1+\epsilon}{2} > \overline{\pi} \geq \frac{1-\frac{n_A}{n_B}\epsilon}{2}$.[25] While the distributional effect clearly increases payments of all

---

[24] We can actually calculate η and ϵ: $\epsilon = \frac{n_A\pi_A+n_B\pi_B}{n_A+n_B} - \frac{(n_A-1)\pi_A+n_B\pi_B}{n_A-1+n_B} = \frac{n_B(\pi_A-\pi_B)}{(n_A+n_B-1)(n_A+n_B)}$ and $\eta = \frac{n_A\pi_A+(n_B-1)\pi_B}{n_A+n_B-1} - \frac{n_A\pi_A+n_B\pi_B}{n_A+n_B} = \frac{n_A(\pi_A-\pi_B)}{(n_A+n_B-1)(n_A+n_B)}$.

[25] As an example, think of the case with $n_A = n_B = 1$ with $\pi_A = 0.99$ and $\pi_B = 0.01$. $\overline{\pi} = 0.5$ in the equivalent homogeneous case. Agents face complete uncertainty about the other agent's type in their incentive compatibility constraints. This decreases individual payments. However, the level of uncertainty is almost zero in the heterogeneous

146

agents, the group composition effect can increase payments of all agents or have different effects on the different groups.

The effect an individual has on the average probability is clearly decreasing in the population size:

$$\frac{\partial \epsilon}{\partial n}, \frac{\partial \eta}{\partial n} < 0$$

$$\lim_{n \to \infty} \epsilon = 0.$$

We can say that the absolute value of the group composition effect is decreasing in population size, whatever the sign of the effect is. This is in contrast to the distributional effect, which is strictly positive and increasing in the population size. It follows that the distributional effect dominates the group composition effect in larger groups.

When discussing the distributional and group composition effects so far, we analysed them under the assumption of identical modes in the distributions. I mentioned this earlier and argued that I think it should not systematically change the results. However, in general, we have to consider how this third effect, which I will refer to as the **modal effect**, will affect individual payments.

Differences in modes are most likely to occur when heterogeneity is high, because the means of the distributions will very often have different integer parts. To provide an intuitive argument of how this affects payments, consider the example of $n_A = n_B = 1$ with $\pi_A = 0.99$ and $\pi_B = 0.01$. There are two possible mechanisms in the two agent case. Either one payment is sufficient for provision or two payments are needed. The mode for agent A is $m_A = 0$, because agent B is very unlikely to value the good. The mode for agent B is $m_B = 1$, because agent A is very likely to value the good. Hence, the optimal cost-splitting rule is not identical for members of different subgroups. We did, however, assume them to be identical so far. Furthermore, agent A is willing to make a very high payment of $P_A = 0.99\theta_H$ in the single payment mechanism, while agent B's maximum payment is very small ($P_B = 0.01\theta_H$). The reverse is true in the cost-splitting mechanism. Not only is the mode very likely to be different when heterogeneity is large, but the pivotality probabilities are such that the minimum payment, which the designer is constrained to when price discrimination is not used, in either mechanism is potentially very small.

The modes being different is independent of the population size. The effect that pivotality probabilities are such that the minimum payment in any mechanism is very small is, however, decreasing in the population size. When I refer to the modal effect, I talk about the combined effect. Even when the modes are identical there is still a modal effect, because the modal probabilities are different.[26]

---

case. Since $\epsilon = \eta = 0.49$, the inequalities hold and uncertainty decreases for both agents, increasing their payments. Obviously, our assumption of identical modes in the distributions is wrong in this case, but we disregard this for the moment.

[26]One might want to separate these effects or argue that the difference in the modal probabilities is actually part

$P_H(l = 0)$ is very different from $P_H(l = 1)$ for both agents in the two agent case. However, the two become much closer as the population size increases; $P_H(l = x) \approx P_H(l = x + 1)$ for large x. Although modes might be different, which can lead to very small payments in small groups, the difference in modes (which is limited to the case of $m_B = m_A + 1$) does not have a big effect on payments in large groups.

We can, therefore, say that the modal effect will decrease payments. The negative effect is increasing in heterogeneity, decreases in absolute terms with population size and converges to zero as $n \to \infty$. Combining the derived effects for the three individual effects yields the overall effect as stated in the above proposition. □

The underlying reason why the group composition effect and the modal effect converge to zero is that they are caused by differences in the incentive compatibility constraints for members of different subgroups. This is simply caused by the fact that each agent i is pooled with $(n_i - 1)$ members of his own subgroup and n_j members of the other subgroup. However, this difference in the incentive compatibility constraints disappears as the population size increases, so that all agents essentially face the same incentive compatibility constraint in very large groups . The distributional effect is not caused by differences in the incentive compatibility constraints, but by the fact that the variation is smaller in the Poisson-Binomial distribution than in the Binomial distribution. This effect persists and, as shown, actually grows with group size.

To summarise, we have established that small heterogeneous groups will underperform equivalent homogeneous groups, both in terms of provision probabilities for given c as well as in terms of maximising EAP, when price discrimination is not used. This is mainly the case because of a negative modal effect, which is caused by not making use of price discrimination, and potentially a group composition effect. Aggregate payments are smaller, because individuals face different incentive compatibility constraints. Some of the constraints are slack without price discrimination. However, both these effects are decreasing in group size, while the positive distributional effect is increasing in group size. Heterogeneous groups will, thus, outperform homogeneous groups in terms of expected aggregate payments as the group size increases. This is due to the fact that overall uncertainty is smaller in the heterogeneous groups. The individual agent's perceived probability of being pivotal is, therefore, larger.

Since maximised expected aggregate payments are larger for heterogeneous groups for sufficiently large n, this mechanism, which, as shown earlier, provides the good in expectation (with probability

of the group composition effect. In any way, these two effects are driven by the same underlying issue; different incentive compatibility constraints for members of different groups.

of $50 + \%$) at the highest possible costs, is able to exist for higher levels of costs in the heterogeneous setting even without making use of price discrimination. The upside of heterogeneity is decreased overall uncertainty increasing pivotality. The downside of not using price discrimination is not making use of different incentive compatibility constraints, but this downside disappears with group size. Since the maximised expected payments are larger in the heterogeneous case, it must be that the budget constraint in the heterogeneous case will be slack at the upper level of costs derived in the homogeneous case. This means that l can be increased and the good provided with a higher probability for given c and n if n is large enough.[27]

In the following, I highlight some of these results using numerical examples in order to convey the magnitudes of the effects. None of the simplifying assumption used in the analysis are used at any point in these numerical examples. Instead, means and modes of the different distributions are calculated. If modes are identical at $m_A = m_B = m^*$, then $P_H = \min[P_{H,A}(m^*), P_{H,B}(m^*)]$. If the modes are different, i.e. $m_A \neq m_B$, then

$P_H = \max\big[\min[P_{H,A}(m_A), P_{H,B}(m_A)], \min[P_{H,A}(m_B), P_{H,B}(m_B)]\big]$. The examples only focus on the behaviour of P_H^*. We are comparing equivalent groups. Whatever the ratio in individual payments is, is also the ratio in expected aggregate payments. The analysis of EAP^* would not give us additional insights.
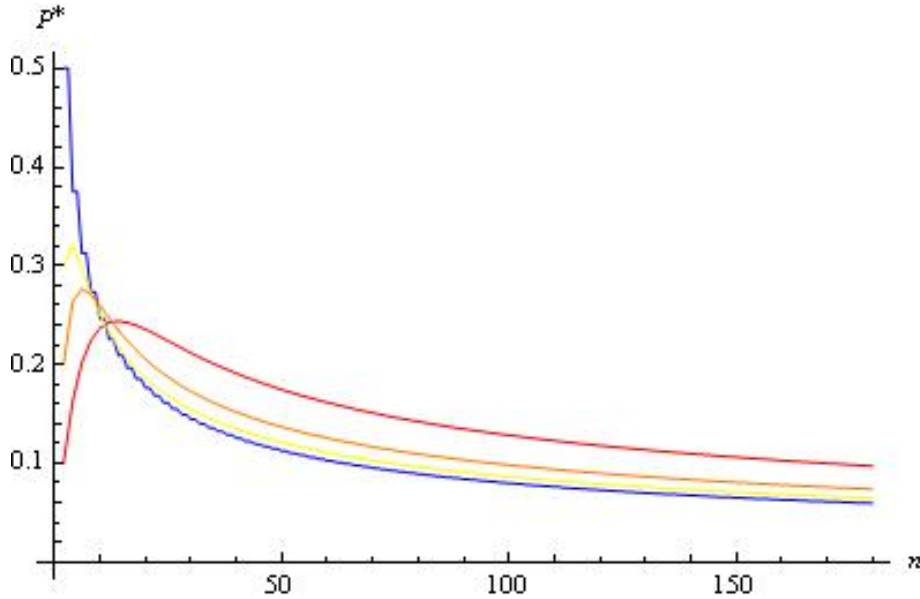

Figure 6: Comparison of individual payments for different levels of heterogeneity

---

[27]Note however, that I have not derived whether provision probabilities for any level of c are higher in the heterogeneus setting in large groups. I have only shown that probabilities around the mode are larger in the Poisson-Binomial distribution, not that this is the case in the tails of the distribution. For example, it might be that the provision probability in homogeneous groups is larger for a very low c, for which a mechanism with high l is feasible, even if the group size is large.

The blue line in figure 6 represents the maximum individual payments in the case of a homogeneous population with $\pi = 0.5$. The yellow line represents the heterogeneous case with $\pi_A = 0.7$ and $\pi_B = 0.3$, orange $\pi_A = 0.8$ and $\pi_B = 0.2$, and red $\pi_A = 0.9$ and $\pi_B = 0.1$. It is assumed that $n_A = n_B$ in all scenarios. As derived in the analytical part, the modal effect is strictly negative and large in small groups. However, the effect is decreasing in the population size, while the distributional effect is increasing. We can see that the combined effect of a falling modal and an increasing distributional effect leads to an actual increase in the maximum payment up to a point where the distributional effect dominates and payments in the heterogeneous case are larger. From there on, the ratio of payments, $\frac{P^*_{Heterogeneous}}{P^*_{Homogeneous}}$, is increasing. In line with the analysis, we can see that the negative effect in small groups caused by the modal effect is increasing in heterogeneity, while the positive effect in large groups, which is solely determined by the distributional effect, is convex in the level of heterogeneity.

The analysis highlighted that the increase in payments is larger, the larger heterogeneity and the more similar the groups are in size. In this sense, the red line in figure 6 can be seen as an upper bound of this increase, because we assumed the groups to be of equal size and, by assuming that $\bar{\pi} = 0.5$, allowed for the maximum possible level of heterogeneity. If we do not assume that $\bar{\pi} = 0.5$, then we cannot have groups of equal size and maximum heterogeneity. We either have to allow for a lower level of heterogeneity or subgroups of different sizes. The next graphs show how this affects payments.
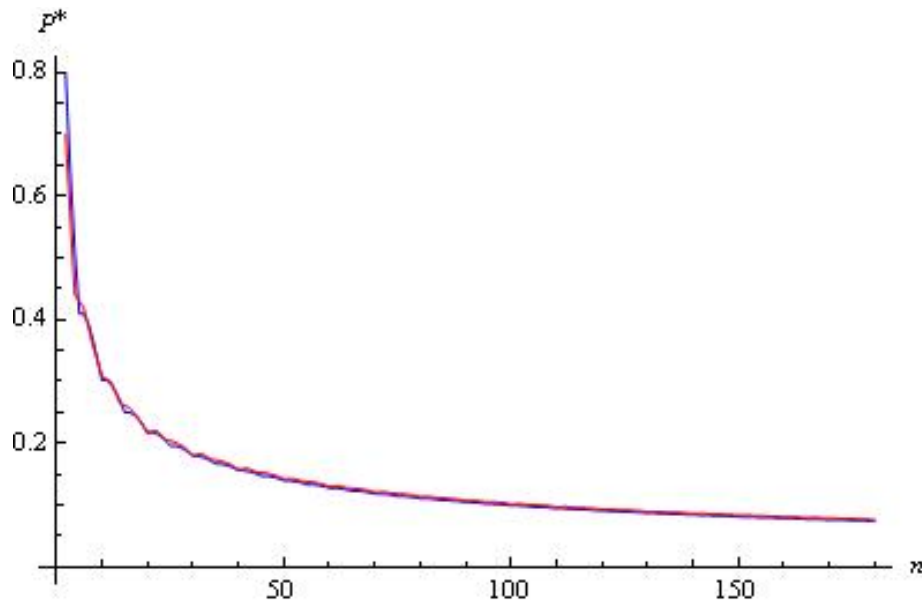

Figure 7: $\bar{\pi} = 0.2$, $\pi_A = 0.3$ and $\pi_B = 0.1$, $n_A = n_B$

In figure 7, it is assumed that $\bar{\pi} = 0.2$ and $n_A = n_B$. This limits heterogeneity, because all probabilities are non-negative. The blue line depicts payments in the homogeneous case, while the red line depicts the case of $\pi_A = 0.3$ and $\pi_B = 0.1$. Again, payments in the heterogeneous case are smaller for small group sizes because of the modal effect, while they are larger for large group sizes. The graph shows that all effects are relatively small, because each effect is increasing in heterogeneity, which is not very high in this example. Note that the case of $\bar{\pi} = 0.8$ would look almost exactly symmetrical, because the only difference would be in the choice of mechanism. A mechanism in which costs are split into more parts would be chosen and the roles of the groups would be inverted. The graph would not be completely symmetrical, because the changes in modes would occur at different points.
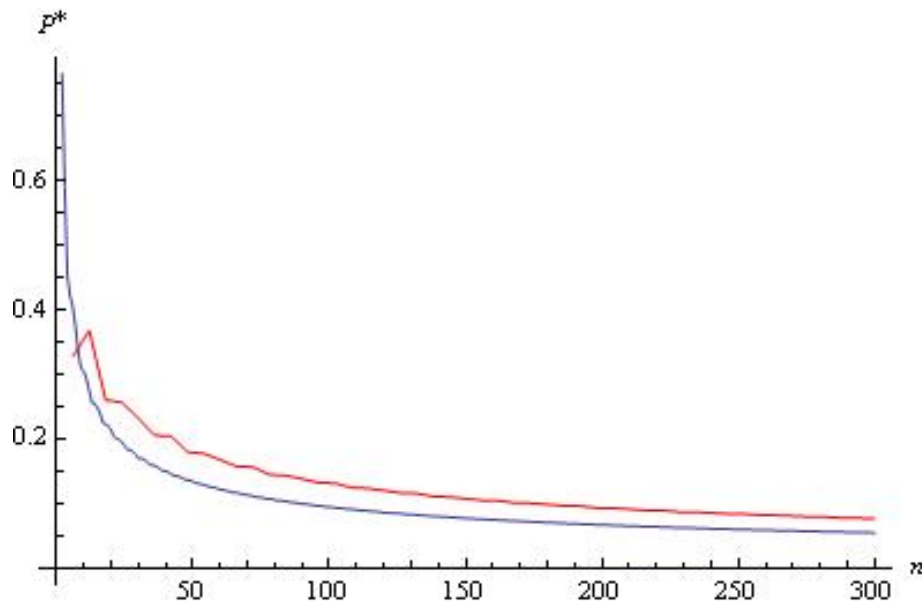

Figure 8: $\bar{\pi} = 0.2\overline{3}, \pi_A = 0.9$ and $\pi_B = 0.1$, $5n_A = n_B$

Figure 8 compares a homogeneous group with $\bar{\pi} = 0.2\overline{3}$, depicted by the blue graph, with the heterogeneous case with $5n_A = n_B$, $\pi_A = 0.9$ and $\pi_B = 0.1$. Subgroups are very unequal in size (in contrast to the previous case), which reduces the increase in payments. On the other hand, I can allow for a high level of heterogeneity by doing this, which magnifies the increase in payments. All effects are now substantially larger in size. The graph is considerably less smooth due to the fact that we can only compare points that are actually comparable. Hence, the red line essentially only has numerical solutions for $n = 6, n = 12, n = 18...$.

2.2.2 Price discrimination

So far, I have analysed the behaviour of payments in the heterogeneous case assuming that only a single payment P_H^* is set. Payments of members of different subgroups were perfect substitutes. I will now allow for price discrimination; $P_{H,A}(\cdot) \neq P_{H,B}(\cdot)$. Agents who have the same valuation ex post can be asked to make different interim payments based on subgroup membership.

In the case of homogeneous agents, or the case without price discrimination, the agents' maximum payments were determined by the probability of being pivotal. This is still the case. However, the probability cannot be simply stated as $f(n-1-l)$. When the payments are non-identical for different agents, agents can be pivotal for different combinations of how many and which agents value the good. When the payments are not identical, it is at least conceivable that agent i could be pivotal independent of whether agent j values the good or not, while agent j could only be pivotal when agent i values the good.

2.2.2.1 Two Agents

I am again going to provide a step-by-step derivation for the two agent case. There are now essentially three different mechanisms to choose from.

In the first case, the designer specifies two payments such that $c \leq \min\left[P_{H,A}(\cdot), P_{H,B}(\cdot)\right]$. This is the equivalent to the *single payment mechanism* from the homogeneous setting. Either payment is large enough for provision and agents are, therefore, willing to contribute up to their total valuation in case they are pivotal. However, the probabilities of being pivotal differ. It follows from the incentive compatibility constraints that individual payments are given by:

$$P_{H,A}(\cdot) = (1 - \pi_B)\theta_H$$

and

$$P_{H,B}(\cdot) = (1 - \pi_A)\theta_H.$$

We can see that $P_{H,A}(\cdot) > P_{H,B}(\cdot)$ and so $\min\left[P_{H,A}(\cdot), P_{H,B}(\cdot)\right] = P_{H,B}(\cdot)$. Agent A is more pivotal in this single payment mechanism, because the other agent is less likely to value the good.

The mechanism produces expected aggregate payments of:

$$EAP(\cdot) = \pi_A(1 - \pi_B)\theta_H + \pi_B(1 - \pi_A)\theta_H$$

a provision probability of:

$$\Pr(\#\theta_H \geq 1) = \pi_A\pi_B + \pi_A(1 - \pi_B) + \pi_B(1 - \pi_A)$$

and satisfies the ex ante budget constraint for:

$$c \le \bar{c}(\cdot) = \frac{\pi_A(1 - \pi_B)\theta_H + \pi_B(1 - \pi_A)\theta_H}{\pi_A\pi_B + \pi_A(1 - \pi_B) + \pi_B(1 - \pi_A)}.$$

However, the initial condition $c \le \min[P_{H,A}(\cdot), P_{H,B}(\cdot)]$ has to be satisfied. Since $1 - \pi_A < \bar{c}(\cdot)$, the mechanism actually only exists for:

$$c \le \bar{c}_{1,PD} = (1 - \pi_A)\theta_H.$$

In the second case, which is going to be a mechanism with *extreme price discrimination*, the costs lie between the smaller and bigger payment: $\min[P_{H,A}(\cdot), P_{H,B}(\cdot)] < c \le \max[P_{H,A}(\cdot), P_{H,B}(\cdot)]$. This scenario could not exist in the homogeneous setting. Theoretically, there are two variations of this mechanism. The higher payment can be set for agent A or agent B.

Suppose $P_{H,B}(\cdot) > P_{H,A}(\cdot)$. It then follows from the incentive compatibility constraints that agent A is never pivotal. He will, therefore, not contribute. Agent B, on the other hand, is always pivotal. It follows that:

$$P_{H,A}(\cdot) = 0$$

and

$$P_{H,B}(\cdot) = \theta_H.$$

The mechanism yields expected aggregate payments of:

$$EAP(\cdot) = \pi_B\theta_H$$

a provision probability of:

$$\Pr(\theta_B = \theta_H) = \pi_B$$

and satisfies the budget constraint for:

$$c \le \bar{c}_{2,PD} = \theta_H.$$

It also satisfies the initial condition put on the mechanism.

Suppose $P_{H,B}(\cdot) < P_{H,A}(\cdot)$. It then follows from the incentive compatibility constraints that agent B is never pivotal, while agent A is always pivotal. Individual payments are given by:

$$P_{H,A}(\cdot) = \theta_H$$

and

$$P_{H,B}(\cdot) = 0.$$

The mechanism yields expected aggregate payments of:

$$EAP = \pi_A\theta_H$$

a provision probability of:

$$\Pr(\theta_A = \theta_H) = \pi_A$$

and satisfies the budget constraint for:

$$c \leq \bar{c}_{2,PD} = \theta_H.$$

It is clear that it is dominant to choose to make agents who are ex ante more likely to value the good pay if one wants to make use of such a mechanism with extreme price discrimination, because $\pi_A \theta_H > \pi_B \theta_H$. Hence, the first of these two options can be ruled out.

Finally, one can specify a *cost-splitting mechanism*, i.e. $\max[P_{H,A}(\cdot), P_{H,B}(\cdot)] < c \leq P_{H,A}(\cdot) + P_{H,B}(\cdot)$. It follows from the incentive compatibility constraints of agents that individual payments under this rule are given by:

$$P_{H,A} = \pi_B \theta_H$$

and

$$P_{H,B} = \pi_A \theta_H.$$

Both payments are needed and agents are, thus, only pivotal if the other agent is a high type. The other agent being a high type is more likely to be the case for agent B. His maximum incentive compatible payment is actually larger. The mechanism yields expected aggregate payments of:

$$EAP(\cdot) = 2\pi_A \pi_B \theta_H$$

a provision probability of:

$$\Pr(\#\theta_H \geq 2) = \pi_A \pi_B$$

and satisfies the budget for:

$$c \leq \bar{c}_{3,PD} = 2\theta_H.$$

Proposition 5.

Allowing for price discrimination to be used in a heterogeneous group with $n_A = n_B = 1$ results in:
- *smaller inefficiencies compared to the same heterogeneous group without the use of price discrimination;*
- *a larger set of costs for which an inefficient mechanism has to be used compared to the equivalent homogeneous group. However, the level of inefficiencies as measured by the difference in provision probabilities is smaller in the heterogeneous case; and*
- *larger expected aggregate payments in the heterogeneous case compared to the equivalent homogeneous group, unless $\pi_A > 0.5$ and $\pi_B > 0.5$, or π_A close to one and π_B close to 0.5.*

While price discrimination clearly increases efficiency levels and the proposition indicates that the heterogeneous two-agent group will, in most scenarios, outperform the equivalent homogeneous group, the proposition shows that there are scenarios in which this is not the case, at least for the two agent case.

Proof. It is obvious that the ability to price discriminate cannot make things worse in terms of efficiency by straightforward logic. One could always opt to set identical payments; the optimal mechanism that does not use price discrimination is still a feasible choice. It is straightforward to verify this further by comparing efficiency measures with and without price discrimination.

What is more interesting is to see how the outcomes compare to the homogeneous case.[28] The designer would like to choose the single payment mechanism, because it has the highest provision probability. However, it only creates small EAP and does, therefore, only satisfy the budget constraint and incentive compatibility constraints up to $c \leq \bar{c}_{1,PD}$, with $\bar{c}_{1,PD} = (1 - \pi_A)\theta_H$. For $c \in [(1 - \pi_A)\theta_H, \theta_H]$, the designer will choose the second mechanism, which is also a single payment mechanism, but one using extreme price discrimination. The designer has to use a cost-splitting mechanism when $c \in [\theta_H, 2\theta_H]$.

Comparing these to the first best results, we can see that inefficiently low provision probabilities arise for $c \in [(1 - \pi_A)\theta_H, \theta_H]$, as was the case in the setting without price discrimination. Price discrimination, however, allows for a mechanism to exist that provides the good with first best probabilities at the upper end of costs.[29] Since $(1 - \pi_A) < \bar{c}(2, 1, \pi)$ for $\pi = \frac{\pi_A + \pi_B}{2}$, the set of costs for which there exist inefficiencies is still larger in the heterogeneous case than in the homogeneous case though. However, because $2\pi(1 - \pi) > \pi_B(1 - \pi_A)$ for $\pi = \frac{\pi_A + \pi_B}{2}$, the level of inefficiencies, measured as the difference between first best provision probabilities and probabilities under the voluntary contribution mechanism, is smaller in the heterogeneous setting. This gives us mixed results in terms of efficiency.

Using our second measure of efficiency, we can see that the heterogeneous case outperforms (creates higher expected payments) in most cases. However, the reverse is true if $\pi_B > 0.5$ (and, therefore, $\pi_A > 0.5$ by definition), or if π_A is close to one and π_B close to a half. The difference is very small whenever the homogeneous case provides larger EAP, while the difference in EAP can be very large in the cases the heterogeneous setting provides larger EAP. We can thus say that price discrimination at least partially offsets the previous inefficiencies caused by not using price

---

[28] Remember that inefficiencies in the two-agent heterogeneous group without price discrimination were considerable worse than in the equivalent two-agent homogeneous group.

[29] Remember that this second set of costs for which there existed no budget satisfying mechanism was an issue when not making use of price discrimination.

discrimination, while general results on efficiency for the two agent case are not clear-cut. □

2.2.2.2 N Agents

In general, l is now not just a number, but at least a vector indicating which agent values the good or not. There could be several vectors indicating all possible combinations under which an agent is pivotal. The mechanism is not described by the number of payments needed for provision, but by combinations of agents.

For example, if $P_{H,A} = 2P_{H,B}$ and agent i is pivotal when x members of group A and x members of group B make a payment, then agent i would also be pivotal if $1.5x$ members of group A and no member of group B make a payment. At the moment, I do not find it feasible to identify a general optimal mechanism analytically.

However, we can still characterise partially such a mechanism by contrasting it to the previous case of no price discrimination. We have seen that two effects had a clear direction in the case of no price discrimination. The modal effect decreased payments, while the distributional effect increased them, with the modal effect dominating in small groups and the distributional effect dominating in larger groups.[30]

The distributional effect always exists (as long as payments are positive for both groups), because it is a purely statistical effect and is not driven by differences in incentive compatibility constraints. The use of price discrimination enables the designer to charge prices according to each individual's incentive compatibility constraint. It will increase payments by not treating them as perfect substitutes. It should then be possible to set payments such that both incentive compatibility constraints are binding.

The two agent example tries to explain the intuition behind price discrimination. What price discrimination can do, is to minimise the modal effect by allowing the designer to set payments according to each agents' modal probability. Payments are not limited by the smallest modal probability. It follows that this effect is potentially very large in small groups, where the modal effect itself is actually large. However, as argued earlier, the incentive compatibility constraints converge to a single incentive compatibility constraint, so that this modal effect, which is addressed by price discrimination, disappears.

There should be no difference between using and not using price discrimination in very large groups; individual payments and aggregate payments are identical. Overall, price discrimination can at

---

[30]As shown, the group composition effect is ambiguous and small in size, especially in larger groups.

least reduce inefficiencies caused by modal effects in small groups and heterogeneous groups tend to be better at providing the public good for large group anyways; price discrimination cannot add anything in large groups.

As an extreme case, there is always the possibility of setting one payment to zero and actually playing the mechanism with a single homogeneous group: the equivalent to the second mechanism in the two agent case. As seen earlier, aggregate payments are strictly increasing in π in the homogeneous case. One would, therefore, always set $P_{H,B} = 0$. This mechanism has the advantage that one only uses the high probability agents and the disadvantage that parts of the population do not make payments at all. In addition to that, the remaining population is homogeneous, so that the positive distributional effect is lost. However, by setting one payment to zero one can extract the maximum payments out of the high probability group.

While I argued that a general optimal price discrimination mechanism should be able to offset the negative modal effect in small groups by setting payments that are not perfect substitutes, but actually reflect the different levels of pivotality, this extreme mechanism gets rid of the modal effect by making the group in effect homogeneous. Obviously, such a mechanism is at least in large groups not the optimal price discrimination mechanism, because the optimal mechanisms using price discrimination must converge to the mechanism not using price discrimination as n increases. I just present this simple mechanism to show that even such a simplistic price discrimination mechanism can at least partially eradicate the inefficiencies that occur in a heterogeneous groups without price discrimination in small groups.
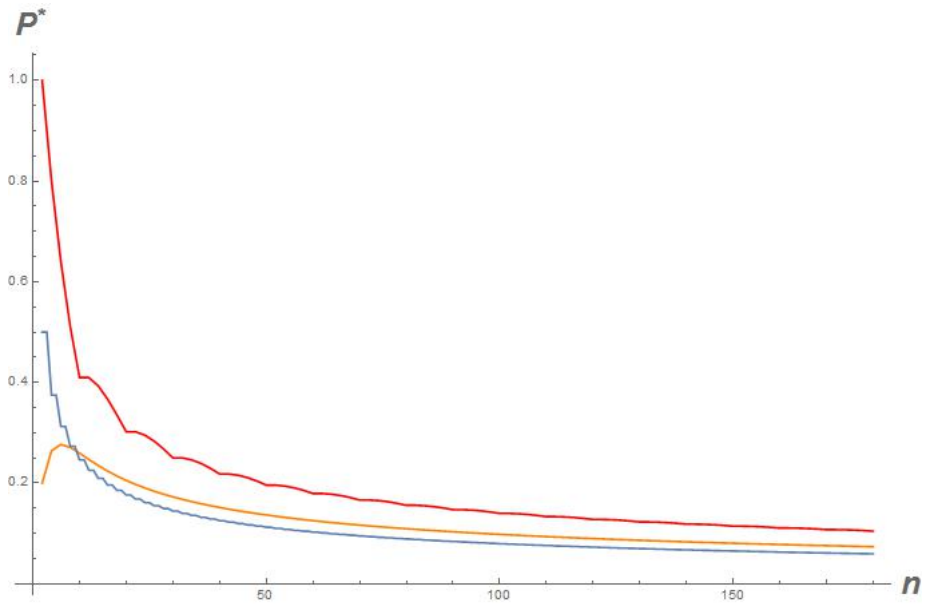
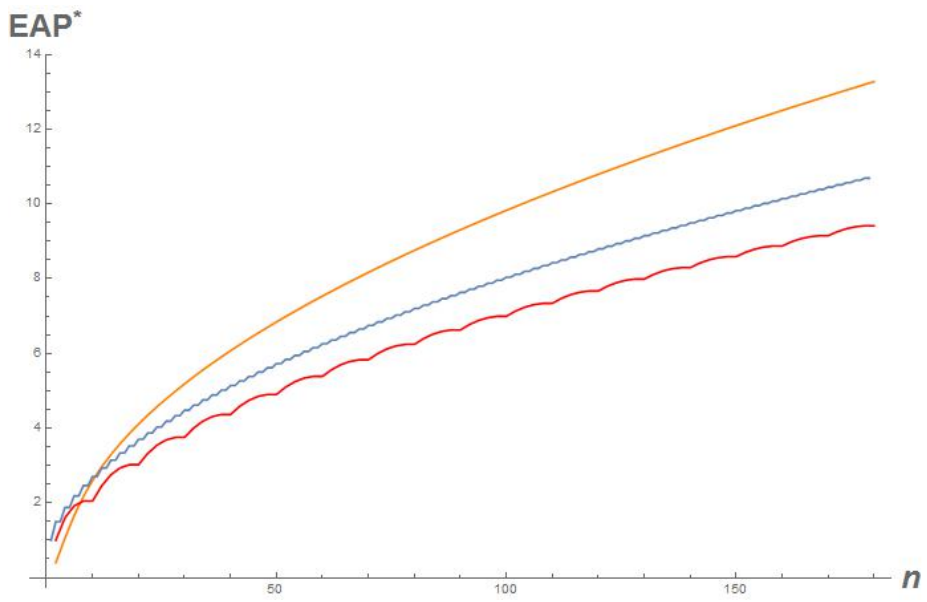Figure 9: P^* for $\overline{\pi} = 0.5$, $n_A = n_B$, $\pi_A = 0.8$ and $\pi_B = 0.2$

Figure 10: EAP^* for $\overline{\pi} = 0.5$, $n_A = n_B$, $\pi_A = 0.8$ and $\pi_B = 0.2$
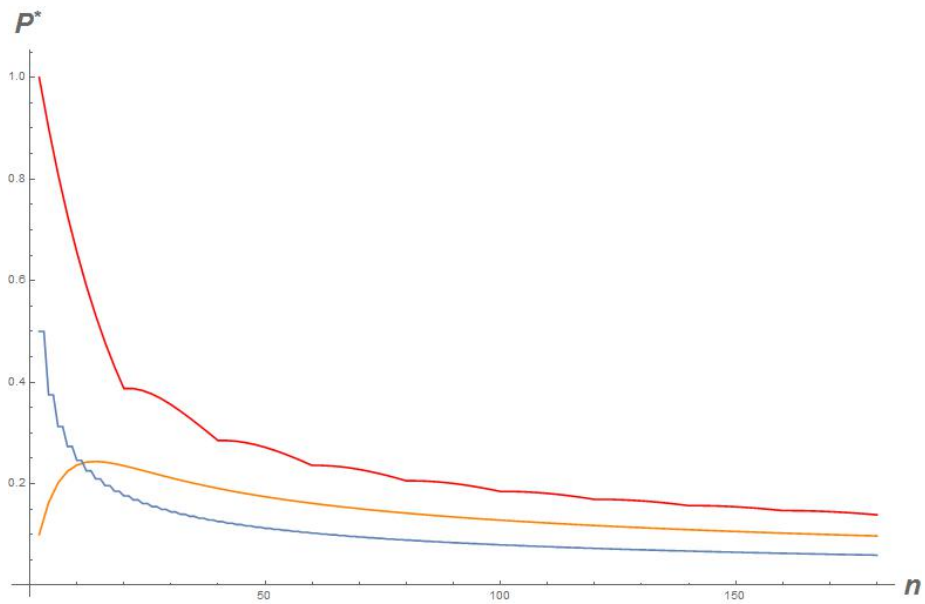
Figure 11: P^* for $\overline{\pi} = 0.5$, $n_A = n_B$, $\pi_A = 0.9$ and $\pi_B = 0.1$
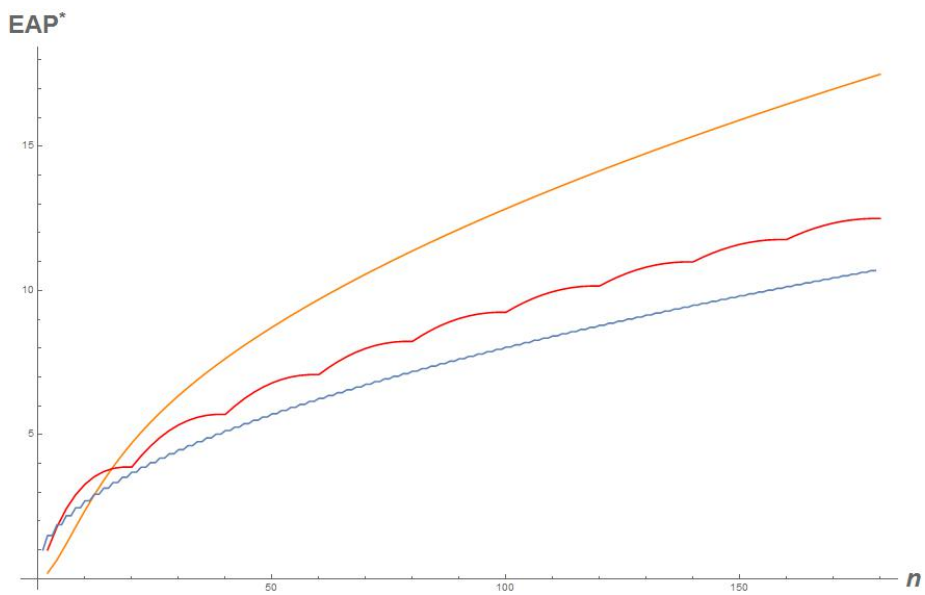

Figure 12: EAP^* for $\overline{\pi} = 0.5$, $n_A = n_B$, $\pi_A = 0.9$ and $\pi_B = 0.1$

The four graphs essentially tell the same story. The red line represents payments for the case of the extreme case of price discrimination, yellow is the case of no price discrimination and blue the equivalent homogeneous group.

The figures show that such an extreme mechanism can increase efficiency compared to the homo-

geneous case when heterogeneity is large. The negative effect caused by a smaller population size is outweighed by the effect of larger individual payments for these cases. In any case, this complete price discrimination mechanism, and in fact any more efficiently chosen price discrimination mechanism, cannot outperform the no price discrimination case in large groups.

3   Conclusion

In a discrete model with a binary public good, interim incentive compatible payments are determined by the probability of being pivotal under a specified mechanism. I have used this to derive (in)efficiency results for small group sizes and have shown how a designer would determine an optimal mechanism for given c and n. Maximum individual payments are determined by the modal probability and are, therefore, decreasing in the population size, while expected aggregate payments are increasing.

Considering the scenario in which the population of agents consists of subpopulations that differ in terms of their average valuations for the public good, I have shown that individual and expected aggregate payments are larger than in the case of a homogeneous group when the group size is large enough. This increase in the level of efficiency is increasing in the level of heterogeneity, both in terms of aggregate payments as well as provision probability for given c.

In general, the designer can make use of price discrimination in the case of a heterogeneous group. Price discrimination allows the designer to make use of the fact that members of different subpopulations face different incentive compatibility constraints. However, this possibility disappears asymptotically as the incentive compatibility constraints converge towards each other. There will, thus, always be a single payment set in large groups. Efficiency levels will, however, still be larger in the heterogeneous group. This increase is due to what I call the distributional effect. Overall uncertainty is smaller in the heterogeneous group, which increases modal probabilities and, hence, payments.

The existence of the modal effect, which effects payments negatively, can severely decrease efficiency in small heterogeneous groups when price discrimination is not used; especially when heterogeneity is large. In this case, incentive compatibility constraints are very different. When he does not make use of price discrimination, the designer restricts himself to the stricter incentive compatibility constraint (leaving the other one slack). While I did not provide a general solution for an optimal price discrimination mechanism in small groups, I provided an intuitive description of such a mechanism. Using a mechanism of extreme price discrimination, I showed that this negative effect can at least

be partially offset. In summary, the tool of price discrimination does not result in an increase in efficiency in large groups. However, in that case, the heterogeneous case outperforms the equivalent homogeneous case anyway. Price discrimination can only be used gainfully in small groups.

References

Al-Najjar, Nabil I.; and Smorodinsky, Rann. "Pivotal Players and the Characterization of Influence", *Journal of Economic Theory*, 2000, 92, pp. 318-342

Boergers, Tilman; and Norman, Peter. "A Note on Budget Balance under Interim Participation Constraints: The Case of Independent Types", *Economic Theory*, 2009, 39(3), pp. 477-489

Cadsby, Charles Bram; and Maynes, Elizabeth. "Voluntary provision of threshold public goods with continuous contributions: experimental evidence", *Journal of Public Economics*, 1999, 71, pp. 53-73

Chan, Kenneth S.; Mestelman, Stuart; Moir, Robert; and Muller, R. Andrew. "Heterogeneity and the Voluntary Provision of Public Goods", *Experimental Economics*, 1999, 2, pp. 5-30

Cherry, Todd L.; Kroll, Stephan; and Shogren, Jason F. "The impact of endowment heterogeneity and origin on public good contributions: evidence from the lab", *Journal of Economic Behavior&Organization*, 2005, 57, pp. 357365

Burlando, Roberto M.; and Guala, Francesco. "Heterogeneous Agents in Public Goods Experiments", *Experimental Economics*, 2005, 8, pp. 35-54

Cramton, Peter; Gibbons, Robert; and Klemperer, Paul. "Dissolving a Partnership Efficiently", *Econometrica*, 1987, 55, pp. 615-632

Darroch, John N. "On the Distribution of the Number of Successes in Independent Trials", *Annals of Mathematical Statistics*, 1964, 35, pp. 1317-1321

d'Aspremont, Claude; and Gérard-Varet, Louis-Andre. "Incentives and Incomplete Information", *Journal of Public Economics*, 1979, 11, pp. 25-45

Fischbacher, Urs; Gaechter, Simon; and Fehr, Ernst. "Are people conditionally cooperative? Evidence from a public goods experiment", *Economics Letters*, 2001, 71, pp. 397404

Fisher, Joseph; Isaac, Mark R.; Schatzberg, Jeffrey W.; and Walker, James M. "Heterogeneous demand for public goods: Behavior in the voluntary contributions mechanism", *Public Choice*, 1995, 85, pp. 249-266

Hellwig, Martin F. "Public-Good Provision with Many Participants", *Review of Economic Studies*, 2003, 70, pp. 589-614

Hoeffding, Wassily. "On the Distribution of the Number of Successes in Independent Trials", *Annals of Mathematical Statistics*, 1956, 27, pp. 713-721

Isaac, R. Mark; Walker, James M.; and Williams, Arlington W. "Group size and the voluntary provision of public goods", *Journal of Public Economics*, 1994, 54, pp. 1-36

Laforgia, A.; and Natalini, P. "Some Inequalities for the Ratio of Gamma Functions", *Journal of Inequalities and Special Functions*, 2011, 2, pp. 16-26

Ledyard, John O.; and Palfrey, Thomas R. "A Characterization of Interim Efficiency with Public Goods", *Econometrica*, 1999, 67(2), pp. 435-448

Ledyard, John O.; and Palfrey, Thomas R. "A general characterization of interim efficient mechanisms for independent linear environments", *Journal of Economic Theory*, 2007, 133, pp. 441-466

Mailath, George J.; and Postlewaite, Andrew. "Asymmetric Information Bargaining Problems with Many Agents", *Review of Economic Studies*, 1990, 57, pp. 351-367

Myerson, Roger B. "Incentive Compatibility and the Bargaining Problem", *Econometrica*, 1979, 47(1), pp. 61-73

Myerson, Roger B.; and Satterthwaite, Mark A. "Efficient Mechanisms for Bilateral Trading", *Journal of Economic Theory*, 1983, 29, pp. 265-281

Myerson, Roger B. "Two-Person Bargaining Problems with Incomplete Information", *Econometrica*, 1984, 52(2), pp. 461-488

Norman, Peter. "Efficient Mechanisms for Public Goods with Use Exclusions", *Review of Economic Studies*, 2004, 71, pp. 1163-1188

Prohorov, Yu V. "The asymptotic behavior of the binomial distribution", *Selected Translations in Mathematical Statistics and Probability*, 1961, 1, pp. 87-95

Rapoport, Amnon; and Ramzi, Suleiman. "Incremental Contribution in Step-Level Public Goods Games with Asymmetric Players", *Organizational Behavior and Human Decision Processes*, 1993, 55, pp.171-194

Rondeau, Daniel; Schulze, William D.; and Poe, Gregory L. "Voluntary revelation of the demand for public goods using a provision point mechanism", *Journal of Public Economics*, 1999, 72, pp. 455470

Samuels, S.M. "On the Number of Successes in Independent Trials", *Annals of Mathematical Statistics*, 1965, 36, pp. 1272-1278

Conclusion and reflections

In my research, I focused on two main areas of interest: the effect of competition in different markets and the voluntary provision of public goods. I extend existing models mainly in two ways: by changing the assumptions about the level of information that agents have and by allowing for an increased level of heterogeneity of agents. This comes at the cost of less tractable models. The set of parameters often becomes too large to solve the model and derive general results.

To counteract this further complication, I needed to take steps that reduced the complexity of the models. I chose to focus on a particular subset of goods in each chapter. In the first chapter, I specifically focus on consumption chains in which a primary good is essential. Restricting myself to this case allows me flexibility in other areas of the model (I can allow for primary and secondary good valuations to differ across agents, and I can change the level of information consumers have). In the second chapter, I focus on reports or information as a good, but assume that it is not sold at a price. This allows me to combine supply and demand side motives into a single model. In the final chapter, I restrict myself to a binary public good and make restricting assumptions on the type of mechanism used. I can analyse the effect ex ante heterogeneity has on efficiency levels of voluntary contribution mechanisms and can analyse efficiency levels in finite group sizes only because I restrict myself to this specific scenario.

The degree to which one needs to make simplifying assumptions in order to maintain tractability in the model and to maintain any hopes of deriving meaningful results surprised me during my research. Introducing increased flexibility, for example by allowing for agent heterogeneity, into existing models is rarely feasible unless one restricts the model in other areas. This forced me to make more restricting assumptions than I had hoped to make at the beginning of each chapter. I had to adjust my expectations of what is feasible. Over time I noticed that it is usually easier to start with a basic model and extend it, rather than to start off with what I considered to be the best representation of the market I wanted to model and then introduce simplifying assumptions until the model could be solved.

I do, though, feel that the existing literature often aims to establish general results that apply as widely as possible. My approach is different. I think building models with the aim to derive eye-catching, general results has the danger that such models often fail to capture the vast complexity of the real world. There are many different types of public goods and many different types of private goods. There is information that is sold and there is information that is distributed freely. Not only is there a vast array of different goods, but there are also many very different agents in terms of their behaviour, their preferences and other characteristics. I do, therefore, feel that

research about the effects of competition or the efficiency of public good provision should be more targeted. General results can be helpful in highlighting the general advantages and disadvantages of competition or the voluntary provision of public goods, but their useful application is usually much more limited. Although I wish that I could have been able to derive more eye-catching and general results, the fact that I did not achieve this in most instances does, therefore, not disappoint me too much. The answer in economics usually is "it depends" and I feel my chapters can highlight some important elements of what "it depends" on and how.