Competition, Conflict and Institutions: Three Essays in Applied Microeconomic Theory

Giovanni Ko

A thesis submitted for the degree of
Doctor of Philosophy in Economics

Department of Economics
London School of Economics and Political Science

January 2012

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others, in which case the extent of any work carried out jointly by me and any other person is clearly identified in it.

I certify that chapter 2 of this thesis was co-authored with Madhav S. Aney. I, Giovanni Ko, contributed over half of the work on this chapter, including the theoretical model and results.

Giovanni Ko
January 2012

Abstract

This thesis consists of three papers on completion and conflict in three distinct but related settings.

The first paper develops a model of tax compliance and enforcement where homogenous agents receive signals about how tolerant the tax authority is of evasion, and where the latter has imperfect means of detecting evasion. The main results show that increasing the quality of the information that taxpayers have about the tax authority's tolerance of evasion may increase compliance. This is because if the signals are sufficiently informative, taxpayers are engaged in Bertrand-like competition: if all taxpayers are evading a similar amount, each will have a strong incentive to evade slightly below that amount in order to escape detection. This logic is directly opposed to the culture of secrecy that prevails in many tax administrations.

The second paper, jointly written with Madhav Aney, deals with the question of how specialists in violence like the military or the police can commit not to abuse their coercive power. The answer that the paper provides is that competition between specialists in violence creates incentives for them not to expropriate from civilians. The main theoretical results are that these incentives become stronger as competition becomes more intense, both in terms of the number of specialists in violence and in the evenness of their strengths. The hypothesis that greater numbers of specialists in violence leads to less expropriation is tested using cross-country regressions and found to be strongly consistent with the data, especially for the case of developing countries.

The third paper analyses the equilibria of two-player imperfectly discriminating contests of the power-form under incomplete information. This paper develops a method for solving for the Bayesian Nash equilibria of such games by working backwards from the equilibrium distributions of effort, rather than forwards from the distributions of the agents' types. This method is used to prove that there exist no distributions of type such that effort is an affine function of the type. The method is used to construct an equilibrium where effort is log-logistically distributed, carrying out comparative statics. This equilibrium is shown to be special in that it exhibits a formal equivalence to that in a contest with complete information.

Contents

List of Figures

List of Tables

Introduction

This thesis consists of three chapters, each of which is a standalone paper, modelling competition in three distinct settings. The first two chapters examine the role of competition in an institutional setting.

Tax authorities are notoriously secretive about their operations, but how justified is this attitude? The first paper addresses this question in the context of a model of tax compliance and enforcement featuring a continuum of taxpayers and a tax authority that can only audit an exogenously determined proportion of them. Taxpayers receive signals about the tax authority's tolerance, i.e., the proportion that will not be audited, and these signals are their only source of heterogeneity. On the other hand, the tax authority receives signals about each taxpayer's evasion and uses this information to select which ones to audit. The main results are that secrecy is indeed better than openness, but only if the accuracy with which the tax authority detects evasion is sufficiently high. Furthermore, when this accuracy is very high, increasing the quality of the information that taxpayers have about the tax authority's tolerance may increase compliance. This is because if the signals of tolerance are sufficiently informative, taxpayers are engaged in Bertrand-like competition: if all taxpayers are evading a similar amount, each will have a strong incentive to evade slightly below that amount in order to escape detection.

The second chapter, jointly written with Madhav Aney, deals with the question of how specialists in violence, like the military or the police, can commit not to abuse their coercive power, i.e., "who guards the guards themselves?" The answer that the paper provides is that "the guards guard each other", i.e., competition between these agents is one of the mechanisms that can deter predation. In our model, even if specialists in violence could expropriate all output costlessly, it is attractive to protect producers from predators. This is because there is a marginal defensive advantage and consequently defence is an effective way to potentially eliminate other specialists in violence, reducing competition and leading to higher future payoffs. Hence, producers can offer transfers to specialists in violence that make defence a dominant strategy, resulting in an equilibrium without predation. We therefore show that internal competition among specialists in violence is enough to keep predatory behaviour at bay and sustain economic incentives even in the absence of threats external to themselves. The main theoretical results are that these incentives become stronger as competition becomes more intense, both in terms of the number of specialists in violence and in the evenness of their strengths. The hypothesis that greater numbers of specialists in violence leads to less expropriation is tested using regressions on a panel of countries and is found to be strongly consistent with the data, especially for the case of developing countries.

Thus, these two chapters show how competition among agents who have an incentive to subvert certain institutional arrangements can improve the functioning of these institutions.

The third chapter is more foundational in nature and examines competition *per se* in the context of two-player imperfectly discriminating contests of the Tullock or lottery type, i.e., contests where the players can increase their probability of winning by increasing their costly effort, but even the player with the highest effort is not guaranteed to win. Contrary to the settings described in the first two chapters, competition here is wasteful and decreases the total surplus, an effect commonly known in the literature as rent dissipation. These contests have been used extensively to model conflict in settings such as wars, litigation and political competition, but almost exclusively under the assumption of complete information about the contestants' types, viz., their valuation of the prize, or equivalently, the cost of effort. This chapter studies such contests in their abstract form, in the case where the contestants' types are private information, about which little is known, especially compared to the case of perfectly discriminating contests such as first-price all-pay auctions. This is done by solving for the Bayesian Nash equilibria of these games by working backwards from the equilibrium distributions of effort, rather than forwards from the distributions of types. This method is used to prove that there exist no equilibria such that effort is an affine function of type. Moreover, this chapter constructs an equilibrium such that efforts are log-logistically distributed. This equilibrium is special in that the equilibrium behaviour is the same as in a game of complete information where each player faces an opponent of the *median* type. Furthermore, the log-logistic equilibrium is shown to have particularly clear and intuitive comparative statics of effort expenditure and rent dissipation. In particular, it is shown that incomplete information results in lower rent dissipation than complete information.

Acknowledgements

I am greatly indebted to my supervisor, Frank Cowell, for his continued support and guidance throughout my years at the LSE. I am very grateful for the patience and respect with which he nurtured my intellectual development.

I am very grateful to Maitreesh Ghatak for his encouragement and helpful comments over the years. Together with my supervisor, he opened my eyes to the beauty of economics.

I am especially indebted to Madhav Aney, with whom I worked on chapter 2 of this thesis, for being a great intellectual partner and a true *kalyāṇa mitta*.

The articles in this thesis have benefited greatly from comments and suggestions from Sanghamitra Bandhyopadhyay, Tim Besley, Konrad Burchardi, Mariano Cena, Bonsoo Koo, Gerard Padró i Miquel, Stefan Penczynski, Miguel Sanchez-Villalba, Oliver Vanden Eynde, Tianxi Wang and seminar participants at LSE, Alicante, NTU and SMU.

I gratefully acknowledge financial support from the LSE. I am particularly indebted to the late Mrs Elisabeth Scitovsky for establishing the Tibor Scitovsky Scholarship and helping students like me follow in the footsteps of her husband.

A special thanks goes to Gudrun Dalibor, David and Daniel Stuckey who have been a second family to me. I would also like to thank my friends Martin Stoynov and Lisha Yan for their help and hospitality in London.

Finally, I would like to thank my parents, my sister and my wife who made this possible through their love and support.

Chapter 1

Signals and secrecy in tax enforcement

1.1 Introduction

Tax agencies are very secretive about their operations and tend to disclose as little information as possible regarding the methods used for selecting taxpayers for audit and the proportion of taxpayers that are audited. For example, in the UK, although the Freedom of Information Act (FOIA) guarantees a right of public access to government information, it "creates an exemption from the right to know if releasing the information would or would be likely to prejudice [...] the assessment or collection of tax, duty, or similar imposition." [1] In the US, where the corresponding FOIA does not make a specific exemption for the Internal Revenue Service (IRS), the latter has been involved in extensive and protracted litigation[2] to try to prevent disclosure of information relating to its workings by appealing to various general exemptions. Tax agencies evidently believe that this secrecy fosters compliance, as the following passage from Roberts v. IRS (1984), quoted in Reinganum and Wilde (1988), shows:

> One of the tools in the arsenal of the IRS which promotes voluntary compliance is the uncertainty in the minds of the taxpayers as to just how much overstepping of the boundaries of strict compliance will bring down the enforcement authority of the agency. The government argues that, by disclosing the guidelines by which it determines which violations are so egregious as to merit enforcement action, it will permit the taxpayer bent on unlawful tax avoidance to conform his conduct not only to the boundary between strict compliance and noncompliance, but also to the boundary between noncompliance which does not merit enforcement action by the IRS, and noncompliance which is so egregious as to prompt the agency to respond.

This passage explains the motivation behind tax agencies withholding information, but are there any incentives for them to do the opposite? As an illustration, suppose that the tax authority has pre-audit information (e.g., third-party reports or statistical estimates) about individual evasion that is sufficiently accurate to allow it to select the most egregious evaders for audit, and suppose further that taxpayers are identical and perfectly informed about how many

[1] See Section 31 of the Freedom of Information Act 2000.

[2] See United States Tax Reporter, P76,556.502, "Freedom of Information Act" for a complete list of such cases.

audits will be carried out. In this scenario, for every possible candidate for the equilibrium level of evasion, each taxpayer has a strong incentive to evade slightly less than that and significantly reduce their chances of being audited. Therefore, prospective evaders are engaged in competition á la Bertrand where they all have an incentive to undercut each other's evasion amount, leading to an equilibrium with full compliance.

This paper explores the consequences of this logic by developing a model where taxpayers receive signals about the tax authority's type, namely its tolerance of evasion, and the tax authority receives an indicator of each taxpayer's evasion, which it uses to select individuals for audit. The main results lend support to the view that secrecy fosters compliance, but only if the accuracy of the tax agency's information about individual evasion is sufficiently high (Proposition 1.11). Even then, we will see that if the tax authority is open enough about how tolerant of evasion it is, i.e., if the taxpayer's signals about its type are sufficiently accurate, more openness leads to more compliance through the aforementioned Bertrand competition effect (Proposition 1.13).

The present paper is closely related to Reinganum and Wilde (1988), who first analysed the question of whether taxpayer uncertainty is beneficial for compliance. Their model features a tax authority that has perfect pre-audit information about taxpayers' evasion, which corresponds to section 1.3.3 here. Their main result is that as taxpayer uncertainty about the tax authority's cost of audit increases compliance first increases and then decreases, but the logic of why this is the case is not discussed. In fact, their model lacks any strategic interaction between taxpayers, so that the competitive effects mentioned above are not present.

Another related paper is Hansen, Krarup, and Russell (2006), who compare the effects on compliance of no information and full information policies, roughly corresponding to section 1.3.1 here, in a regulatory context. The present paper is more general as it allows for individuals to be imperfectly informed about the number of audits the authority carries out, not just being completely informed or uninformed. Another generalisation here is allowing the authority to have some pre-audit information about individuals on which it can base its audit strategy. Lastly, in thir model individuals are faced with a binary decision of whether to comply or not, whereas here we will consider continuous compliance decisions.

More broadly, the contribution of this paper to the literature on tax compliance is in incorporating signals of the tax authority's type and taxpayers' actions. The importance of the latter is that it enables the tax authority to select taxpayers for audit based on the potential for additional revenue, rather than relying on random audits. Such endogenous audit rules result in a coordination game between taxpayers, which has been analysed in a laboratory setting by Alm and McKee (2004) and Tan and Yim (2011). Indeed, the empirical findings in the latter contradict the view that secrecy is conducive to compliance.

On the theoretical side, one paper that incorporates tax agency information about taxpayer actions is Macho-Stadler and Perez-Castrillo (2002), but in their case the signal is of taxpayers' income, not evasion. This is in contrast to the most notable example of statistical information used by tax authorities, namely the IRS's DIF score: one of the very few publicly divulged facts about the DIF is that it specifically measures the likelihood of evasion. Also, the present model allows for continuous taxpayer actions, rather than the three allowed in Macho-Stadler and Perez-Castrillo (2002).

There is also a sizeable literature on endogenous audit rules *with pre-commitment*, i.e., when the tax authority announces such a rule and commits to it, as in Reinganum and Wilde (1985), but this literature is fundamentally at odds with the notion of secrecy, since the audit rule must be publicly declared, so that the question addressed in this paper cannot be examined in that context.

The presence of signals received by taxpayers about the tax authority's type makes the model in this paper a *global game*, connecting it to the literature on such games summarised in Morris and Shin (2003). Within that literature, the present paper is most closely related to the paper on tax evasion by Sanchez-Villalba (2006). The fundamental difference between that paper and the present one is that in the former, the tax authority chooses whom to audit based on their income report rather than on pre-audit information about evasion itself. Even though taxpayers can in principle evade continuous amount, in equilibrium they either evade or comply fully, with the tax authority's audit rule being similarly binary in nature, in contrast with the equilibria of the present paper where there is a continuous distribution of evasion across taxpayers.

Finally, a more methodological innovation of the current paper is that it considers a signal of the tax authority's type that is not the sum of the true type and an independent error, as is generally the case in the literature on global games and in Sanchez-Villalba (2006), in particular. As we will see in section 1.2.3, the non-additive structure of the signal offers certain advantages over the additive one, especially when the underlying type is distributed over a bounded interval.

The paper is structured as follows: section 1.2 explains the model and its assumptions; section 1.3 establishes the model's equilibria and their properties for select values of the key parameters; section 1.4 analyses the comparative statics of the model, providing the main results of the paper; section 1.5 discusses these results, pointing out their implications and limitations; section 1.6 provides concluding remarks and directions for further research. Proofs and results of a purely mathematical nature are relegated to appendix 1.A.

1.2 Model

We represent the fact that there are many taxpayers by modelling them as a continuum. In particular, we will denote the set of taxpayers by I and identify it with the unit interval $[0, 1]$. With this assumption we capture the notion of negligibility of individuals and normalise their mass to 1.

The tax authority has a private type called *tolerance*, denoted by the random variable T that takes values in $[0, 1]$: T is the mass of taxpayers it lets off so that $1 - T$ are audited and, if found to have evaded taxes, punished. The distribution of T is assumed to be common knowledge. After a particular realisation t of T, taxpayers receive a signal $S^t : I \to [0, 1]$ of t, where $S^t(i)$ is taxpayer i's signal. The relationship between that tax authority's tolerance and the signal thereof is common knowledge and is captured by the *openness* parameter ρ described in detail in section 1.2.3.

Upon learning his signal, each taxpayer updates his beliefs about the value of t and, simultaneously with all other taxpayers, chooses the amount of tax to evade in order to maximise expected utility. Taxpayers have identical preferences that are further spelt out in section 1.2.1. We will use the function $X : I \to \mathbb{R}$ denote the amount[3] of evasion across the population of taxpayers, with $X(i)$ being the amount evaded by taxpayer $i \in I$.

After taxpayers make their choices, a signal $Y : I \to \mathbb{R}$ is generated, with $Y(i)$ being the *indicator* of taxpayer i's evasion $X(i)$. The process by which Y is generated from X is common knowledge and is described in detail in section 1.2.2 where we introduce the *accuracy* α of the indicator. The tax authority uses this indicator to select which taxpayers to audit. The tax authority's objective is to maximise the total amount of evaded taxes collected after audits, together with associated penalties. Upon audit, the full amount of evasion is perfectly discovered and a taxpayer who is found to have evaded x has to pay $(1 + p)x$, where $p > 0$ is a proportional[4] penalty.

To summarise, the timing is:

1. Tax authority learns the realised value t of its type.
2. Taxpayers receive signals S^t of t.
3. Taxpayers choose evasion simultaneously, forming X.
4. Indicators Y of evasion are generated from true evasion X.
5. Tax authority observes Y and chooses whom to audit.
6. Evaded taxes and penalties are collected.

For notational convenience, we will treat the functions S, X and Y as random variables[5] and we will use probabilistic terminology like the cumulative distribution function $F_X(x) := \Pr(X \le x)$ of X to express, for example, the mass of taxpayers who evade at most x. Similarly, we will write $X \overset{d}{=} Y$ to mean that X and Y have the same distribution.

1.2.1 Taxpayers' decision

Taxpayers, which could be individuals or firms, choose how much tax to evade, where evasion consists of anything that illegally reduces tax payments, ranging from underreporting income to claiming tax deductions to which one is not entitled. Taxpayers are risk neutral so they maximise expected wealth, or equivalently, minimise expected payments to the tax authority. If a taxpayer with true tax bill τ, who evades an amount x, is audited with probability π, then his expected payment is $(1 - \pi)(\tau - x) + \pi(\tau + px) = \tau - x(1 - (1 + p)\pi)$, which is equivalent to maximising $x(1 - (1 + p)\pi)$. Dividing throughout by $1 + p$, the objective of the taxpayer is to maximise

$$u(x) = x\big(1/(1 + p) - \pi(x)\big)\,, \tag{1.1}$$

[3]This notation only deals with pure strategies, but we will limit discussion of mixed strategies to section 1.3.1, where this notation will not cause problems.

[4]This is standard in the literature and is true for many tax administrations, including the UK and US, where p is in the order of 20%.

[5]More formally, we can do this by endowing the set I with measure-theoretic structure through the σ-algebra $\mathcal{C}$ of Lebesgue-measurable subsets of I, and the Lebesque measure μ, so that we have the standard probability space $(I, \mathcal{C}, \mu)$. Then we would interpret $\Pr(X \le x)$ as merely shorthand for $\mu\big(\{i : X(i) \le x\}\big)$ or $\mu\big(X^{-1}([0, x])\big)$.

with respect to x. Note that in general, π will depend on the taxpayer's own evasion x, all taxpayers' evasion X, the tax authority's audit strategy and the signal of the tolerance type.

We assume that taxpayers are perfectly informed about their true tax liability and therefore never overpay; upon audit, overpayment does not attract any "bonus" that compensates for the probability that it might not be discovered at all, so x must always be non-negative. Undoubtedly, the assumption that taxpayers cannot make mistakes is rather unrealistic in the face of the complexities of real-life taxation, but we can nonetheless shed some light on the question posed in this paper without relaxing it, so we will not do so, in line with the rest of the literature on the subject.

We will assume that there is an upper bound to the extent of evasion possible, and more importantly, that this upper bound is common to all taxpayers, and is normalised to 1, resulting in the following assumption.

Assumption 1.1 Taxpayers can choose evasion x in $[0, 1]$.

This assumption implies that all taxpayers have the same bounded evasion opportunities. One might object to this on the grounds that if everyone owes the same amount τ of taxes, then anybody who pays less than τ will be automatically detected as an evader. To get around this, we can assume that the true tax bill is unknown to the tax authority prior to an actual audit, because, for example, incomes or profits vary year by year, as in Sanchez-Villalba (2006). For this to be plausible the population of taxpayers needs to be restricted to a fairly narrow audit class, such as self-employed individuals, or firms in a particular sector. In that case, another way to justify assumption 1.1 is to assume that incomes or profits are homogeneous but taxpayers differ in terms of eligible deductions, like business expenses or mortgage interest; maximum evasion would then correspond to claiming enough deductions to reduce taxable income to zero.

Alternatively, we can think of x as the evaded amount as a proportion of the true tax liability, with the proviso that taxpayers cannot evade more than what was due. Likewise, we would have to assume that the tax authority does not consider the worst evaders to be those who have evaded the greatest absolute amount, but those who have evaded the greatest proportion of their tax bill. Support for this view comes from the observation that if tax agencies cared solely about absolute amounts, they would not waste resources auditing "small fish" who cannot be expected to yield large amounts of additional revenue, something that actually does occur.

Whichever way we interpret assumption 1.1, its effect is to ensure that differences in the taxpayers strategies are due solely to differences in the information they receive. This makes the model a global game[6], but we will defer discussing this connection to section 1.5.

1.2.2 Information about evasion

The tax authority selects which taxpayers to audit based on the indicator Y. We can think of this as a synthesis of all information available to the tax authority regarding an individual's

[6] A global game is a coordination game where agents receive a signal about a parameter that affects their payoffs. See Morris and Shin (2003) for an overview.

tax compliance prior to an actual audit, such as third-party reporting (i.e., information from employers, suppliers or financial institutions) and statistical techniques, the leading example of which is the IRS's Discriminant Index Function (DIF).[7] With a loose enough interpretation, this information could even include such intangible factors as "gut feelings" that a tax inspector has about a tax return, based perhaps on past experience. The upshot of this is that the information embodied in Y is in general not verifiable, or at least insufficient to collect additional revenue or impose any penalties, for which an actual audit is required, perhaps even by law.

Concretely, we will assume that for each taxpayer $i \in I$, the indicator $Y(i)$ gives the correct amount $X(i)$ of evasion with probability α and a randomly drawn value $X(j)$ from the overall distribution of evasion with probability $1-\alpha$. As an example of the role of α, if all taxpayers are complying, an individual who deviates and evades a positive amount will be detected with probability α and will blend in with the others with probability $1-\alpha$. In light of this, we will call α the *accuracy* of Y in measuring X. Treating X and Y as random variables, we can define the process more formally as follows.

Assumption 1.2 The indicator Y is generated from X so that

$$\Pr(Y \leq y \mid X = x) = \begin{cases} (1-\alpha)F_X(y) & \text{if } y < x \\ (1-\alpha)F_X(y) + \alpha & \text{if } y \geq x, \end{cases} \tag{1.2}$$

where $F_X(x)$ is the cumulative distribution function of X.

Intuitively, this means that the distribution of Y given $X = x$ has a probability mass of α at $Y = x$, and is otherwise the same as the distribution of X but scaled by the remaining probability $1-\alpha$. The marginal distribution of X and the conditional distribution of Y fully determine the joint distribution of X and Y, which turns out to be symmetric, as stated in the following lemma.

Lemma 1.1 Assumption 1.2 implies that (a) the joint distribution of X and Y is symmetric, and that (b) the correlation between X and Y is α.

PROOF. See proof 1.A.1 in appendix 1.A. □

Lemma 1.1 has a number of implications for the information carried by Y. Firstly, the fact that α is the correlation between X and Y confirms our intuition that α measures the accuracy with which Y reflects X. Secondly, the symmetry of X and Y means that $X \overset{d}{=} Y$, so that by observing Y, the tax authority learns the true distribution of evasion X. Furthermore, we will show in the proof of proposition 1.1 that Y is strongly informative about X in that the posterior

[7] Jones (2000–2001) provides the following description: "DIF uses a mathematical formula for each class of return to measure the probability of error on the return. For example, if a taxpayer's tax return shows more deductions than normal for the taxpayer's income, then DIF will select the return for examination. The actual operation of DIF is highly complex and is a closely guarded secret within the IRS. Basically, the way DIF works is that various items on a taxpayer's tax return are compared with information contained in the computer. Tax returns are, in effect, given a grade and those returns receiving a failing grade are selected for examination. Thus, all returns are scrutinized in the same uniform way. Each return is scored relative to other returns of the same type by a formula and, there, automatically classified and assigned or not assigned for examination according to these relative scores."

distribution of X conditional on a higher value of Y first-order stochastically dominates[8] that for a lower value, meaning that for any hypothetical evasion level x, the higher the indicator, the greater the probability that the taxpayer evaded at least x. This in turn implies that the conditional expectation of X given $Y = y$ is increasing in y.

1.2.3 Information about tolerance

The tax authority's type is its tolerance of evasion, which is the proportion of taxpayers it will not audit, i.e., the complement of the audit rate. This tolerance is a function of budgetary and other constraints that the tax agency faces every time period, such as the amount of funding it receives and the allocation of staff to audit versus non-audit roles.[9] Taxpayers do not have full knowledge of these constraints and how they determine audit rates, so they assess the tax authority's tolerance using all kinds of information, ranging from public announcements by the tax agency itself to word-of-mouth or even guesswork. The more open the tax authority is, for example by releasing credible information about the resources it has available to audit taxpayers, the closer the their estimates are to the true value of tolerance. Conversely, the more secretive it is, the greater the role of myths and hearsay[10] in shaping taxpayers' assessments of their probability of being audited.

We model this by representing the tax authority's type, viz., its tolerance, by a random variable T. The taxpayers' assessment of this is represented by signals S^t they receive after a particular realisation t of T, where $S^t(i)$ is the signal received by taxpayer $i \in I$. As mentioned earlier, we will treat S^t as a random variable. Furthermore, we define $S := S_T$ to be the random variable[11] denoting the *ex-ante* (before T is realised) values of the signal. This enables us to write $S^t \equiv S \mid T = t$ and posit a joint distribution[12] of S and T. We can then write $T^s := T \mid S = s$ to denote the random variable representing the assessment of T by a taxpayer who receives a signal $S = s$.

[8] For more details on this see Milgrom (1981), who was the first to propose first-order stochastic dominance as a characterisation of informativeness of signals.

[9] For example, GAO (2001) notes that: "According to IRS officials, audit rates declined for fiscal years 1996 to 2000 for three main reasons. First, over this period, the number of IRS auditors for individual returns declined by more than half for reasons such as a decline in total staff and decisions to change staffing priorities to better serve taxpayers before they file their returns. Second, IRS was more likely to use the remaining auditors in other duties, such as assisting taxpayers. Third, audits took longer due to additional audit requirements, such as more written communications with taxpayers about the status of their audit."

[10] In the US, there is a cottage industry that churns out books with titles like "What the IRS Doesn't Want You to Know" and "How to Beat the I.R.S. at Its Own Game: Strategies to Avoid – and Fight – an Audit". Given the vigour with which the IRS engages in litigation to prevent publication of information about its internal procedures (see footnote 2), the value of such books is dubious at best.

[11] Formally, if Ω is the sample space of T, so that $T : \Omega \to [0, 1]$, then $S_T : I \times \Omega \to [0, 1], (i, \omega) \mapsto S_{T(\omega)}(i)$.

[12] It is tempting to model the signal for a given t by assuming that there is a continuum of i.i.d. random variables S_i^t, one for each taxpayer i, and then appealing to a suitable law of large numbers to assert that the realised values of the signal across the population of taxpayers have the same distribution as that of S_i^t. If the number of agents were countable, this would indeed be the most natural approach and would pose no problems, but Judd (1985) and Feldman and Gillies (1985) showed that for a continuum of random variables, the required law of large numbers does not exist in the standard measure-theoretic paradigm. Frameworks where such laws of large numbers do exist have been proposed, most notably by Sun (1998) and Sun (2006), but they remain largely inaccessible for non-specialists due to the reliance on non-standard analysis. Since we are not concerned with independence of individual signals per se, we can instead consider the much simpler scenario where the process that generates the signals *somehow makes sure* that they are distributed across the population according to the conditional distribution of S given $T = t$. We shall not delve any deeper into the question of how exactly such a process would work, and simply trust its existence.

We will assume that the joint distribution of S and T is such that it has a parameter ρ that captures the degree of stochastic dependence[13] between S and T. In particular, without loss of generality, we will assume that ρ ranges from 0 to 1, with $\rho = 0$ representing independence and $\rho = 1$ representing perfect dependence or comonotonicity[14]. Moreover, we require that the distribution be continuous in ρ. We will take ρ to be the *openness* of the tax agency, with $\rho = 0$ corresponding to the case of complete secrecy and $\rho = 1$ to the case where all taxpayers know the value of the tax authority's tolerance.

In order to capture the idea that these signals are in some sense "correct", we will make the following assumption.

Assumption 1.3 The joint distribution of tax authority's tolerance T and its signal S is symmetric.

The reason why this assumption means the signals are correct is that it implies that $S \overset{d}{=} T$, i.e., that the marginal distributions of T and S are identical, which in turn implies that $\mathbb{E}S = \mathbb{E}T$. Note that we cannot require that the signal be interim-unbiased, i.e., that it be conditionally unbiased for a given realisation of T so that $\mathbb{E}(S^t) = t$ for all values of t and ρ. As reasonable as that may sound, this requirement is incompatible with the assumption that the joint distribution is continuous in ρ, because as $\rho \to 0$, the joing distribution of S and T tends towards independence, so that $\mathbb{E}(S^t)$ tends to $\mathbb{E}(S) \neq t$.

In fact, the requirement that we be able to accommodate independence between the distributions of the type and its signal rules out the commonly used assumption that the signal be the sum of the type and an independent error term. If the type is distributed over a bounded interval, with a zero-mean additive error, the signal will always convey *some* information about the type, since values of the signal in the support of the distribution of the type are more likely than those outside.

As a consequence of assumption 1.3, the distributions of $S|T$ and of $T|S$ are symmetric, so we will denote the cdf of both by $F(\cdot\,|\,\cdot)$, so that $F(t|s) = \Pr(T \leq t \mid S = s)$. Similarly, we will denote the cdf of S and T by $F(\cdot)$, so that $F(t) = \Pr(T \leq t)$.

To characterise the informativeness of S as a signal of T, we will make the following assumption about the posterior distribution of T.

Assumption 1.4 If $s_1 > s_0$, the posterior distribution of T conditional on $S = s_1$ dominates that of T conditional on $S = s_0$ in monotone likelihood ratio (MLR).

Monotone likelihood ratio dominance is widely used in statistics, and more relevantly to our purposes, in the study of monotone comparative statics. If $F(t|s)$ admits a density[15] $f(t|s)$, assumption 1.4 means that for $t_1 > t_0$, the likelihood ratio $f(t_1|s)/f(t_0|s)$ is increasing in s, i.e., that the relative likelihood of higher values of the type increases with higher values of the signal. Note that monotone likelihood ratio dominance is a relatively strong requirement, certainly

[13] Although the notation ρ is suggestive of mere correlation, it is meant to capture true stochastic dependence. We will see in section 1.3.3 a concrete example of how to model this dependence structure. See in particular footnote 22.

[14] Formally, S and T are comonotonic if there exists an increasing function ϕ such that $S = \phi(T)$.

[15] See Definition 1 in Ormiston and Schlee (1993) for a general, but extremely unwieldy definition of MLR dominance that accommodates distributions that are not continuous.

stronger than first-order stochastic dominance, but it cannot be relaxed for the purposes of this paper, in particular, for proposition 1.2.

Finally, we will make the assumption about average tolerance.

Assumption 1.5 The distribution of tolerance satisfies $1 - \mathbb{E}T < \frac{1}{1+p}$.

Recall that $1 - T$, the complement of tolerance, is the proportion of taxpayers who are audited. If auditing is random and taxpayers' signals are completely uninformative then $1 - \mathbb{E}T$ is the exogenous probability of audited, as in the classical case of Allingham and Sandmo (1972). We see from eq. (1.1) that in this scenario, a risk-neutral taxpayer evade 0 if this probability exceeds $1/(1 + p)$ and will evade 1 otherwise. The assumption that $1 - \mathbb{E}T < 1/(1 + p)$ rules out the uninteresting scenario where the tax authority has such abundant resources that it can ensure full compliance even in the classical case with risk-neutrality.

1.3 Equilibrium

Having laid out the model, we now turn to analysing its equilibria, and we start by describing the optimal audit strategy of the tax authority. The equilibrium concept we use is Bayesian Nash equilibrium, so that the taxpayers' strategies are functions that map signals to evasion. On the other hand, in general, the tax authority's strategy is a function that maps each realisation of its type and the indicator Y to the set of taxpayers to be audited. Fortunately, as the following proposition shows, we need not consider strategies in such generality as the optimal audit strategy is of a particularly simple form.

Proposition 1.1 Suppose $\alpha > 0$ and that the tax authority's tolerance is t. Given evasion indicator Y, let θ be the quantile function of Y, defined by $\theta(t) = \inf\{y : \Pr(Y \leq y) \geq t\}$. Then the best response of the tax authority is to audit all taxpayers $i \in I$ with $Y(i) > \theta(t)$, none with $Y(i) < \theta(t)$ and some with $Y(i) = \theta(t)$, if any.

PROOF. From lemma 1.1 we know that the distribution of X and Y is symmetric, so that

$$\Pr(X \leq x \mid Y = y) = \begin{cases} (1-\alpha)F_X(x) & \text{if } x < y \\ (1-\alpha)F_X(x) + \alpha & \text{if } x \geq y. \end{cases} \tag{1.3}$$

Then, given two taxpayers with indicators y_0 and y_1, where $y_0 < y_1$, the cumulative distribution functions of $X \mid Y = y_0$ and $X \mid Y = y_1$ coincide for $x < y_0$ and $x \geq y_1$, but for $y_0 \leq x < y_1$, $\Pr(X \leq x \mid Y = y_0) = (1-\alpha)F(x) + \alpha > (1-\alpha)F(x) = \Pr(X \leq x \mid Y = y_1)$, since $\alpha > 0$. Hence $\Pr(X \leq x \mid Y = y_0) \geq \Pr(X \leq x \mid Y = y_1)$ for all x, which means that $X \mid Y = y_1$ first-order stochastically dominates $X \mid Y = y_0$, i.e., $\Pr(X \mid Y = y_1 > x) \geq \Pr(X \mid Y = y_0 > x)$ for all $x \in [0, 1]$. Therefore, auditing someone with $Y = y_1$ will bring in at least as much additional revenue as auditing someone with $Y = y_0$, so it is optimal for the tax authority to audit the top $1 - t$ of taxpayers and let off the bottom t, sorted by Y. The threshold $\theta(t)$ for audit is therefore the t-th quantile of Y, i.e., $\theta(t) = \inf\{y : \Pr(Y \leq y) \geq t\}$. If $\Pr(Y = \theta(t)) = 0$, there is no mass at $\theta(t)$ so that exactly t taxpayers are below $\theta(t)$ and $1 - t$ are above, so having an indicator value greater than or equal to $\theta(t)$ will indeed trigger an audit. On the other hand, if $\Pr(Y = \theta(t)) > 0$, $\Pr(Y < \theta(t)) \leq t$ so $t - \Pr(Y < \theta(t)) \geq 0$ of those with $Y = \theta(t)$ are not audited, leaving $\Pr(Y = \theta(t)) - (t - \Pr(Y < \theta(t))) = \Pr(Y \leq \theta(t)) - t$ to be audited randomly. □

This shows that if $\alpha > 0$, i.e., the indicator is not completely uninformative, the tax authority will employ a cutoff strategy[16]. Because of the binary nature of the tax authority's optimal strategy, we will find it convenient to refer to it simply by the threshold $\theta(t)$. Similarly, for convenience, we will note the following fact.

Corollary 1.1 The audit threshold $\theta(t)$ is non-decreasing in t.

Proof. This follows immediately from its definition as the quantile function of $Y \overset{d}{=} X$. □

Having characterised the optimal audit strategy of the tax authority, we now establish some basic properties of the best responses of the taxpayers.

Proposition 1.2 Optimal evasion is non-decreasing in the signal s.

Proof. Given the audit strategy θ, a taxpayer with signal s faces a random audit threshold $\theta(T^s)$ with cdf $F_\Theta(\theta|s)$ and chooses x to maximise

$$U(x,s) = \int_0^1 u(x,\theta)\,dF_\Theta(\theta|s) = \int_0^1 x(1/(1+p) - P(x,\theta))\,dF_\Theta(\theta|s) \tag{1.4}$$

where $P(x,\theta)$ is the probability of being audited when the threshold is θ and evasion is x, which, by proposition 1.1 equals 1 if $x > \theta$, 0 if $x < \theta$, and some number in $[0,1]$, the precise value of which is irrelevant, that depends on the evasion of other taxpayers if $\theta = x$. We can then apply a powerful result of Athey (2002, Theorem 2 together with extension (ii) to Lemma 5) that implies that if $u(x,\theta)$ satisfies the single-crossing property with respect to x, θ and s orders $F_\Theta(\theta|s)$ by monotone likelihood ratio (MLR), then the maximiser of $U(x,s)$ is non-decreasing in s, so that in order to prove the result, we need to check that the conditions on $u(x,\theta)$ and $F_\Theta(\theta|s)$ are satisfied.

Firstly, $u(x,\theta)$ satisfies the single-crossing condition with respect to x, θ if, for all $x_1 > x_0$ and $\theta_1 > \theta_0$, $u(x_1,\theta_0) \geq u(x_0,\theta_0)$ implies $u(x_1,\theta_1) \geq u(x_0,\theta_1)$, with the second inequality being strict if the first one is. If $x_1 > \theta_0$, then $P(x_1,\theta_0) = 1 > 1/(1+p)$, so $u(x_1,\theta_0) < 0$. This is clearly less than $u(x_0,\theta_0)$ if the latter is non-negative, but this is also true if $u(x_0,\theta_0)$ too is negative, because $u(x_1,\theta_0) = -p/(1+p)x_1 < -p/(1+p)x_0 = u(x_0,\theta_0)$, since $x_1 > x_0$. Therefore, $u(x_1,\theta_0) \geq u(x_0,\theta_0)$ implies $x_1 \leq \theta_0$, which in turn implies that $x_0 < x_1 \leq \theta_0 < \theta_1$. Hence $P(x_1,\theta_1) = 0$ and $P(x_0,\theta_1) = 0$, so that $u(x_1,\theta_1) = x_1/(1+p) > x_0/(1+p) = u(x_0,\theta_1)$, as required. Figure 1.1 illustrates this argument.

Secondly, recall that MLR order is preserved by non-decreasing transformations, so that if T^{s_1} dominates T^{s_0} in MLR, then $\theta(T^{s_1})$ dominates $\theta(T^{s_0})$ in MLR. Therefore, by assumption 1.4 and corollary 1.1, the family of cdfs $F_\Theta(\theta|s)$ are ordered by s in MLR. □

[16] This is at first glance similar to what it would do if it could announce and commit to an auditing strategy before evasion decisions are taken, as in Reinganum and Wilde (1985) and Sanchez and Sobel (1993), but the similarity is only superficial because in those models the cutoff rule is conditional on the income report of the taxpayer, since that is the only piece of information available to the authority.

Figure 1.1: Single-crossing for $u(x, \theta)$

Note that $u(x, \theta)$ can take any value in $[-p\theta/(1+p), \theta/(1+p)]$ when $x = \theta$ without affecting this condition.

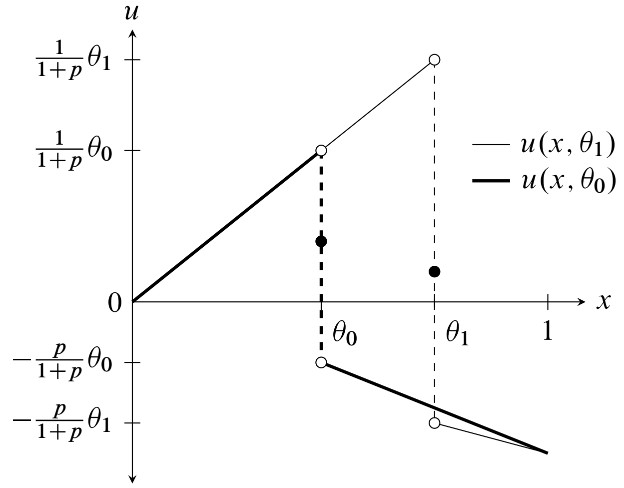

Next, if we call equilibria where all taxpayers evade the same amount *pooling* equilibria, borrowing the term from adverse selection[17], we can state the following proposition about them.

Proposition 1.3 If $\alpha > 0$, the only pooling equilibrium consists of full compliance by all taxpayers.

PROOF. Suppose that there is another pooling equilibrium where all taxpayers evade the same amount x^*. Then, the audit threshold is $\theta(t) = x^*$ for all t, so the payoff of a taxpayer with signal s is $u^* = x^*\big(1/(1+p) - (1 - \mathbb{E}T^s)\big)$. If this taxpayer deviates to evading $\tilde{x} < x^*$, he is audited only when the indicator takes value x^* with probability $1 - \alpha$, so that his payoff is now $\tilde{u} = \tilde{x}\big(1/(1+p) - (1-\alpha)(1 - \mathbb{E}T^s)\big)$. Any taxpayer with s such that $\mathbb{E}T^s < 1$ can choose $\tilde{x}$ sufficiently close to x^* to ensure that $\tilde{u} > u^*$, thus yielding a strictly profitable deviation. The only value of x^* where this cannot occur is $x^* = 0$, so this must be the only pooling equilibrium. □

The intuition behind this result is as follows. When all taxpayers are evading the same amount, each taxpayer has an incentive to slightly undercut everyone else and thus significantly reduce his audit probability, leading to an equilibrium with no evasion. Thus, if the parameters are such that in equilibrium taxpayers act homogeneously, they are engaged in a Bertrand-like game with the familiar "race to the bottom" effect mentioned in the introduction.

The next proposition establishes the conditions under which this zero evasion equilibrium exists.

Proposition 1.4 Full compliance is an equilibrium if $\alpha(1 + p) \geq 1$. Moreover, full compliance

[17]The present model is not an example of adverse selection because although taxpayers' signals S^t are private information, they do not affect the tax authority's payoff.

is an equilibrium if and only if, for $\alpha < 1$,

$$\mathbb{E}T^s \leq \frac{1 - 1/(1+p)}{1-\alpha} \quad \text{for all } s \in \operatorname{supp}(S), \tag{1.5}$$

where $\operatorname{supp}(S)$ is the support of S.

Proof. Suppose now that all taxpayers evade 0, so that their payoff is 0 and the audit threshold is $\theta(t) = 0$ for all t. If a taxpayer with signal s deviates to $x > 0$, his indicator will take value x with probability α and 0 with probability $1 - \alpha$, so that he will be audited with probability $\alpha + (1-\alpha)(1 - \mathbb{E}T^s) \geq \alpha$. Then his payoff from this deviation will be $x\big(1/(1+p) - \alpha - (1-\alpha)(1-\mathbb{E}T^s)\big)$, which is at most the equilibrium payoff 0 iff $1/(1+p) - \alpha - (1-\alpha)(1-\mathbb{E}T^s) \leq 0$, which is equivalent to eq. (1.5).

If $\alpha(1+p) \geq 1$, $\frac{1-1/(1+p)}{1-\alpha} \geq 1 \geq \mathbb{E}T^s$ for any s, since $\operatorname{supp}(T) \subseteq [0,1]$, so that eq. (1.5) is satisfied.

Finally, if $\alpha = 1$, any deviation from 0 will be caught with probability 1, yielding a payoff of $x(1/(1+p) - 1) < 0$, so that full compliance is indeed an equilibrium. $\square$

This means that zero evasion can be an equilibrium if accuracy and penalties are sufficiently high, as intuition would suggest.

Having established these general results, we will examine the four extreme values of the openness and accuracy parameters, viz., $\alpha = 0$, $\rho = 0$, $\rho = 1$, and $\alpha = 1$.

1.3.1 Full secrecy and openness

Consider the case $\rho = 0$, where the signal S is independent of type of which it is a signal, T. We refer to this case as that of *full secrecy*, since the taxpayers have no information about the particular realisation of the tax authority's tolerance and make their evasion decisions based solely on their commonly known prior information. In this case, the unique pure strategy[18] equilibrium is as follows.

Proposition 1.5 For $\rho = 0$, if $\alpha < \frac{1/(1+p)-(1-\mathbb{E}T)}{1-(1-\mathbb{E}T)}$, in the unique pure strategy equilibrium, taxpayers are indifferent between evading any $x \in [\underline{x}, 1]$ and the distribution of evasion across taxpayers is given by the cdf

$$F_X(x) = \begin{cases} 0 & \text{if } x < \underline{x} \\ F^{-1}\left(\frac{1}{\beta}\left(1 - \frac{1-\beta}{x}\right)\right) & \text{if } x \in [\underline{x}, 1] \\ 1 & \text{if } x > 1, \end{cases} \tag{1.6}$$

[18]We express the equilibria in this section in terms of asymmetric pure strategies, when perhaps it would be more natural to express them in terms of symmetric mixed strategies. Intuitively, since there is a continuum of agents, we would expect these two to be observationally equivalent by appealing to some suitable law of large numbers, so that there is no loss of generality in considering pure strategies only. But it turns out that problems similar to those alluded to in footnote 12 plague this reasoning. As a consequence, for the overall distribution across agents to be the same as that followed by each individual through some idiosyncratic randomisation device, it must be that these randomisation devices are correlated with each other in such a way that they jointly induce the required population distribution. Therefore the symmetric equivalent of the pure strategy equilibria discussed in this section are not mixed strategy ones (where the randomisation is independent across individuals) but more general *correlated* equilibria.

where $\underline{x} = 1 - \beta$ and $\beta := \alpha / \big(1/(1+p) - (1-\alpha)(1-\mathbb{E}T)\big)$. If $\alpha \geq \frac{1/(1+p)-(1-\mathbb{E}T)}{1-(1-\mathbb{E}T)}$, there is no evasion.

Proof. Since all taxpayers have the same information, in order for them to evade different amounts it must be that they are indifferent between them. Let $\underline{x}$ denote the minimum amount of evasion and note that the maximum amount of evasion must be 1: if it were some $\bar{x} < 1$, then evading 1 would yield the same probability of audit as evading $\bar{x}$, since there are no taxpayers who evade more than $\bar{x}$, so that the payoff is strictly higher, resulting in a profitable deviation.

Let F_X be cdf of the distribution of evasion across the population. Assuming F_X is continuous and increasing, the tax authority's audit threshold will satisfy $\Pr(X \leq \theta(t)) = t = F_X(\theta(t))$, so that $\theta(t) = F_X^{-1}(t)$. Then the probability of being audited when evading x is

$$\pi(x) = \alpha \Pr(x \geq \theta(T)) + (1-\alpha)(1-\mathbb{E}T) = \alpha \Pr(x \geq F_X^{-1}(T)) + (1-\alpha)(1-\mathbb{E}T) \tag{1.7}$$

$$= \alpha \Pr(F_X(x) \geq T) + (1-\alpha)(1-\mathbb{E}T) = \alpha F(F_X(x)) + (1-\alpha)(1-\mathbb{E}T)\,. \tag{1.8}$$

Then, for taxpayers to be indifferent between any evasion amount $x \in [\underline{x}, 1]$, it must be that for some $\bar{u} \in [0, 1]$,

$$u(x) = x\,(1/(1+p) - \alpha F(F_X(x)) - (1-\alpha)(1-\mathbb{E}T))) = \bar{u} \quad \text{for all } x \in [\underline{x}, 1] \tag{1.9}$$

$$\iff F_X(x) = F^{-1}\left(\frac{1}{\alpha}\left(\frac{1}{1+p} - (1-\alpha)(1-\mathbb{E}T) - \frac{\bar{u}}{x}\right)\right)\,. \tag{1.10}$$

Furthermore, since $\underline{x}$ and 1 are the minimum and maximum evasion levels, $F_X(\underline{x}) = 0$ and $F_X(1) = 1$, so that

$$F_X(1) = 1 \iff \bar{u} = 1/(1+p) - (1-\alpha)(1-\mathbb{E}T) - \alpha \tag{1.11}$$

$$F_X(\underline{x}) = 0 \iff \underline{x} = 1 - \frac{\alpha}{1/(1+p) - (1-\alpha)(1-\mathbb{E}T)} \tag{1.12}$$

so that, letting $\beta := \frac{\alpha}{1/(1+p)-(1-\alpha)(1-\mathbb{E}T)}$,

$$F_X(x) = F^{-1}\left(\frac{1}{\beta}\left(1 - \frac{1-\beta}{x}\right)\right)\,. \tag{1.13}$$

For this to be a valid equilibrium we also need $\underline{x} \in [0, 1]$, so we need to check when these conditions are satisfied. Firstly, $\underline{x} \leq 1$ iff $\beta \geq 0$ iff $1/(1+p) \geq (1-\alpha)(1-\mathbb{E}T)$, which is always satisfied, since by assumption 1.5, $1/(1+p) > 1 - \mathbb{E}T \geq (1-\alpha)(1-\mathbb{E}T)$. On the other hand, $\underline{x} \geq 0$ iff $\beta \leq 1$ iff $\alpha \leq \frac{1/(1+p)-(1-\mathbb{E}T)}{1-(1-\mathbb{E}T)} < 1$, so accuracy must be sufficiently low.

Note that as $\alpha \to \frac{1/(1+p)-(1-\mathbb{E}T)}{1-(1-\mathbb{E}T)}$, $\beta \to 1$, so that $F_X(x) \to 1$ for all x, and the equilibrium becomes the full compliance one. Indeed, from proposition 1.4, this is an equilibrium iff $\mathbb{E}T \leq \frac{1-1/(1+p)}{1-\alpha} \iff \alpha \geq \frac{1/(1+p)-(1-\mathbb{E}T)}{1-(1-\mathbb{E}T)}$. □

Consider next the case of $\rho = 1$, where taxpayers' signals are perfectly accurate. We refer to this case as *full openness* since the taxpayers have complete knowledge of the tax authority's tolerance, resulting in a game of complete information. The next result describes the (essentially) unique pure strategy equilibrium in this case.

Proposition 1.6 For $\rho = 1$, if $t > \frac{1-1/(1+p)}{1-\alpha}$, t taxpayers evade x^* and $1-t$ taxpayers evade 1, where $x^* := 1 - \alpha/\big(1/(1+p) - (1-\alpha)(1-t)\big)$ is the audit threshold. If $t = \frac{1-1/(1+p)}{1-\alpha}$, at least t taxpayers evade 0 and the rest evade 1. If $t < \frac{1-1/(1+p)}{1-\alpha}$, there is no evasion.

Proof. Recall that $1-t$ is the number of audits that the tax authority can carry out. Since t is known and by proposition 1.1 the tax authority always follows a cutoff strategy, in equilibrium the audit threshold is common knowledge.

Suppose that the equilibrium threshold is some $\theta^* > 0$, so that the probability of being audited when evading x is $(1-\alpha)(1-t)$ if $x < \theta^*$ and $\alpha + (1-\alpha)(1-t)$ if $x > \theta^*$, and some intermediate value if $x = \theta^*$, where $(1-\alpha)(1-t)$ is the probability of being audited when the indicator is incorrect.

If $t > \frac{1-1/(1+p)}{1-\alpha}$, which is only possible if $1/(1+p) > \alpha$, then it is not optimal to evade $x \in (\theta^*, 1)$ because evading 1 has the same probability of audit but with a higher payoff when not audited, since $1/(1+p) - \alpha - (1-\alpha)(1-t) > 0$. Furthermore, for the equilibrium threshold to be θ^*, at most $1-t$ taxpayers must be evading 1 and some must be evading $x = \theta^*$, so let $x^* := \theta^*$. For this to be the case, taxpayers must be indifferent between evading x^* and 1, which means that $u(x^*) = x^*(1/(1+p) - (1-\alpha)(1-t)) = 1/(1+p) - 1 - (1-\alpha)(1-t) = u(1)$ which holds iff $x^* = 1 - \alpha/\big(1/(1+p) - (1-a)(1-t)\big)$. Note that if $t > \frac{1-1/(1+p)}{1-\alpha}$, $x^* = \theta^* > 0$, as supposed, and also $x^* \leq 1$. Moreover, exactly $1-t$ taxpayers must be evading 1: if not, then the probability of being audited when evading x^* would be greater than $(1-\alpha)(1-t)$, so that deviating to $x < \theta^* = x^*$ would yield a strictly higher payoff for x sufficiently close to x^*.

If $t = \frac{1-1/(1+p)}{1-\alpha}$, which is only possible if $1/(1+p) \geq \alpha$, $1/(1+p) - \alpha - (1-\alpha)(1-t) = 0$, so evading $x > \theta^*$ yields a payoff of 0. Suppose that $\theta^* > 0$, then for θ^* to be an audit threshold it must be that some taxpayers are evading $x = \theta^*$, which implies that $1-t$ taxpayers are evading $x > \theta^*$ by the same reasoning as in the last paragraph. But then, taxpayers must be indifferent between evading θ^* and evading $x > \theta^*$, and since the latter yields a payoff of 0, whilst the former a payoff of $\theta^*(1/(1+p) - \alpha)$, it must be that $\theta^* = 0$, a contradiction. Hence, θ^* must be 0, so that there must be at least t taxpayers complying, with the rest evading any positive amount.

If $t < \frac{1-1/(1+p)}{1-\alpha}$, by proposition 1.4 full compliance is an equilibrium. It is also the only equilibrium, because for any audit threshold θ^*, nobody will evade $x > \theta^*$ as that yields a negative payoff, since $1/(1+p) - \alpha - (1-\alpha)(1-t) < 0$. Hence the only possibility is that everyone evades 0 and the audit threshold is also 0. □

The equilibrium is unique except for the case of $t = \frac{1-1/(1+p)}{1-\alpha}$, where there are infinitely many equilibria, spanning all possibilities between the limits of the equilibria in the other two cases, so that the equilibrium varies continuously in t.

1.3.2 Random auditing

Consider the case $\alpha = 0$. Here, the indicator conveys no information about how much tax each individual has evaded, so the tax authority must pick taxpayers at random for audit, rather than using a cutoff rule. The probability of audit is therefore independent of evasion, so that taxpayers will evade 0 or 1, depending on what their assessment of the tolerance of the tax authority.

Proposition 1.7 If $\alpha = 0$, taxpayers with a signal s such that $1 - \mathbb{E}T^s < 1/(1+p)$ evade 1, those with $1 - \mathbb{E}T^s > 1/(1+p)$ evade 0, and those with $1 - \mathbb{E}T^s = 1/(1+p)$ are indifferent between any evasion amount.

PROOF. This follows immediately from the fact that the payoff of a taxpayer who receives a signal s is $u_s(x) = x\big(1/(1+p) - (1 - \mathbb{E}T^s)\big)$, which is increasing, constant or decreasing in x depending on whether $1/(1+p)$ is greater than, equal to or less than $1 - \mathbb{E}T^s$. □

By assumption 1.5, when $\rho = 0$ and taxpayers decide how much to evade based on their prior about T, there will be full evasion. On the other hand, when $\rho = 1$, there will be full evasion whenever $1 - t < 1/(1+p) \iff t > p/(1+p)$ and full compliance when $t < p/(1+p)$, so that ex-ante evasion is less than when $\rho = 0$.

1.3.3 Perfect indicator

We now turn to the case where $\alpha = 1$, i.e., the tax authority has perfect knowledge of a taxpayer's evasion before an actual audit[19] Here we will make the following specific assumption about the joint distribution of the type T and signal S.

Assumption 1.6 Let $\tilde{S}$ and $\tilde{T}$ have the standard bivariate normal distribution with correlation coefficient ρ. Then $S := \Phi\big(\sqrt{1-\lambda^2}\tilde{S} + \lambda\Phi^{-1}(\mu)\big)$ and $T := \Phi\big(\sqrt{1-\lambda^2}\tilde{T} + \lambda\Phi^{-1}(\mu)\big)$, where Φ is the standard normal cdf, $\lambda \in [0, 1)$ and $\mu \in (0, 1)$.

Intuitively, T and S are weighted[20] averages of a normal random variables and a constant, transformed by the normal cdf to yield values in the unit interval. The weighting parameter λ captures the concentration of the prior distribution of tolerance: as $\lambda \to 1$, the distribution become more and more concentrated around μ, whereas as $\lambda \to 0$, the distribution tends towards the uniform distribution on the unit interval[21], commonly used as an ignorance prior. The fact that $\tilde{S}$ and $\tilde{T}$ are standard bivariate normal implies that S and T are symmetric

[19] Recall that an audit is still necessary to collect the evaded tax because the pre-audit information that the tax authority has is assumed to be non-verifiable.

[20] Note that the *squares* of the weights add up to 1. This merely a cosmetic assumption that makes the prior and posterior distributions similar in form and easier to manipulate algebraically. Making the weights themselves add up to 1 does not make any qualitative difference to the results.

[21] Recall that if X has cdf F, then $\Pr(F(X) \le u) = \Pr(X \le F^{-1}(u)) = F(F^{-1}(u)) = u$, which is the cdf of the uniform distribution, so that $F(X) \sim U[0, 1]$.

Figure 1.2: Plots of the density $f(t|s)$

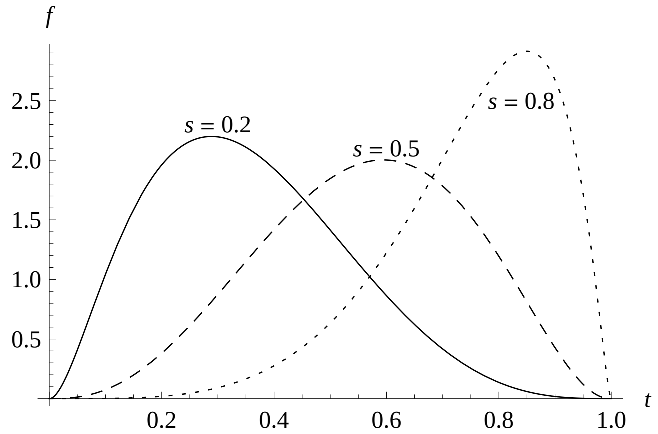

(a) $\rho = 0.7, \lambda = 0.7, \mu = 0.8$

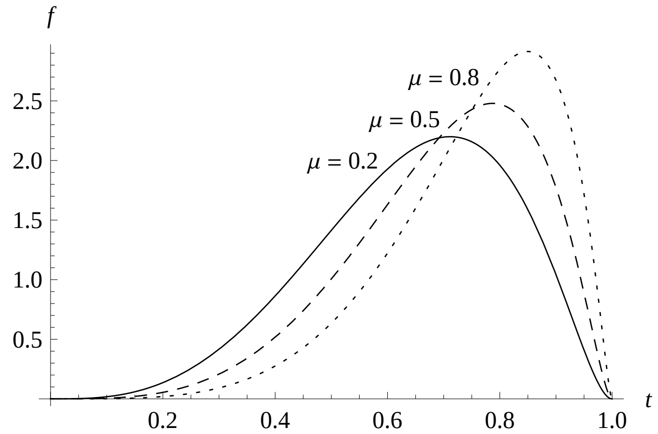

(b) $s = 0.8, \rho = 0.7, \lambda = 0.7$

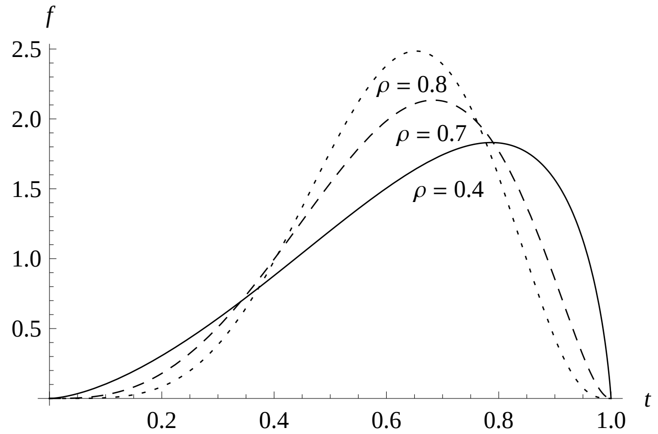

(c) $s = 0.6, \lambda = 0.7, \lambda = 0.8$

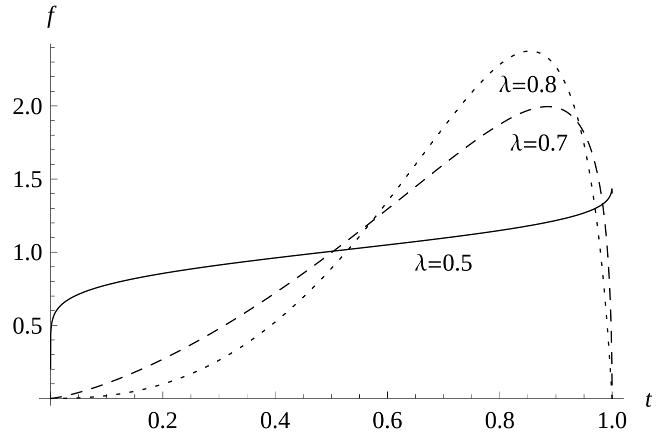

(d) $s = 0.8, \rho = 0, \mu = 0.8$

and that ρ measures the degree of dependence[22] between the two, with $\rho = 0$ meaning they are independent and $\rho = 1$ meaning that $S = T$. Moreover, the joint normality of $\tilde{S}$ and $\tilde{T}$ means that conditional distributions $S|T$ and $T|S$ are of an analytically tractable form, as lemma 1.A.1 shows. Lastly, lemma 1.A.2 shows that assumption 1.6 implies that assumption 1.4 is satisfied.

We begin working towards the equilibria of the model, by setting down the following useful definitions and result.

Lemma 1.2 Let $F^{-1}(s|t)$ be the inverse of $F(s|t)$ with respect to s and let $G(t) := F^{-1}(t|t)$. Then the inverse $G^{-1}(t)$ of $G(t)$ exists for all $t \in [0, 1]$.

[22] Recall that if X and Y are bivariate normal, then their correlation coefficient captures the degree of dependence between the two. Another way of expressing assumption 1.6 is to say that the dependence structure of S and T is captured by the bivariate normal copula $C(\cdot\,, \cdot\,; \rho)$ so that the joint cdf of S and T is $C(F(s), F(t); \rho)$. C could in principle be any copula that includes the independence and comonotonicity ones, but we will use the Gaussian one because of analytical convenience. Copulas are now widely used in finance and risk management to model the dependence structure of random variables. See Nelsen (2006) for the theory and McNeil, Frey, and Embrechts (2005) for applications.

PROOF. See proof 1.A.2 in appendix 1.A □

Since $F(s|t)$ is the cdf of the distribution of signals conditional on the true type of the tax authority being t, $F^{-1}(s|t)$ is the s-th quantile of the distribution of signals, i.e., a mass s of taxpayers will receive a signal less than or equal to $F^{-1}(s|t)$, so that $G(t) := F^{-1}(t|t)$ is the threshold value of the signal such the mass of taxpayers who receive a signal below this threshold is exactly equal to the tax authority's tolerance. Since by proposition 1.2 evasion is non-decreasing in the taxpayer signal, we will see that when the tax authority has tolerance t, it will choose its audit threshold so that it audits those taxpayers who received a signal above $G(t)$, letting off all those who received a signal below $G(t)$. Hence, $G^{-1}(t)$ is the value of the tax authority's tolerance such that in equilibrium precisely those who receive a signal below t are let off and those with a signal above t are audited. Conversely, in equilibrium, a taxpayer who receives a signal t is audited if the tax authority's tolerance is below $G^{-1}(t)$ and is let off if it is above $G^{-1}(t)$.

We will restrict our attention to symmetric pure strategy equilibria where the taxpayers' strategies consist of the same function $x^*(s)$ mapping the signal s into evasion. The first result describes how the tax authority's equilibrium strategy depends on the taxpayers'.

Lemma 1.3 The audit threshold is $\theta(t) = x^*(G(t))$.

PROOF. Let x^{*-1} be the quasi-inverse of x^* defined by $x^{*-1}(x) := \sup\{s : x^*(s) \leq x\}$, equal to the ordinary inverse if x^* is increasing rather than merely non-decreasing, with the property that $\Pr(Z \leq x^{*-1}(x)) = \Pr(x^*(Z) \leq x)$ for any random variable Z with support $[0, 1]$. Then

$$\theta(t) = \inf\{x : \Pr(x^*(S^t) \leq x) \geq t\} = \inf\{x : \Pr(S^t \leq x^{*-1}(x)) \geq t\} \tag{1.14}$$
$$= \inf\{x : F(x^{*-1}(x)|t) \geq t\} = \inf\{x : x^{*-1}(x) \geq F^{-1}(t|t)\} \tag{1.15}$$
$$= \inf\{x : x \geq x^*\big(F^{-1}(t|t)\big)\} = x^*\big(F^{-1}(t|t)\big) = x^*\big(G(t)\big), \tag{1.16}$$

as required. □

Next, we show that taxpayers cannot pool at any level of evasion other than 0.

Lemma 1.4 If there exist $s_0 < s_1 \in [0, 1]$ such that $x^*(s) = c$ for all $s \in [s_0, s_1]$, then $c = 0$.

PROOF. Let s_0 and s_1 be the smallest and largest values in $[0, 1]$ such that x^* takes constant value c for all $s \in [s_0, s_1]$. Then, by lemma 1.3, $G(t) \in [s_0, s_1]$ for all $t \in [G^{-1}(s_0), G^{-1}(s_1)]$, where $G^{-1}s$ exists for any $s \in [0, 1]$ by lemma 1.2, so that $\theta(t) = x^*(G(t)) = c$. This means that for a taxpayer who received a signal s, the probability of being audited when evading c is $\Pr(T^s < G^{-1}(s_0)) + \Pr(G^{-1}(s_0) \leq T^s \leq G^{-1}(s_1))$. But for someone who receives a signal $s \in [s_0, s_1]$, c is not the optimal amount of evasion, because by evading $c' < c$ instead, the taxpayer faces a probability of audit of only $\Pr(T^s \leq G^{-1}(x^{*-1}(c'))) < \Pr(T^s < G^{-1}(s_0))$, since s_0 is the lowest s for which $x^*(s) = c$ and x^* is non-decreasing, thus avoiding audit with probability $\Pr(G^{-1}(s_0) \leq T^s \leq G^{-1}(s_1)) > 0$ since $s_0 < s_1$. Therefore, for c' sufficiently close to c, the payoff from this deviation is strictly greater than that from c, so that c cannot be the equilibrium level of evasion for someone receiving signal $s \in [s_0, s_1]$. The only value of c for which this argument doesn't hold is $c = 0$, so that if x^* is constant over an interval it must be 0. □

This means that the only constant section of x^* must be where there is no evasion, i.e., the distribution of evasion can have a mass point only at $x = 0$, being continuous everywhere else. Moreover, by lemma 1.3 and lemma 1.2, this means that there exists a $t_0 \in [0, 1]$ such that $\theta(t) = 0$ for all $t \leq t_0$ and such that $\theta(t)$ is strictly increasing for all $t > t_0$.

As the final step before computing the equilibrium, we need the following definitions and result.

Lemma 1.5 Let $H(t|s) := F\big(G^{-1}(t)|s\big)$ and let

$$H(s) := H(s|s) = \Phi\left(\frac{-k\Phi^{-1}(s) - \lambda\sqrt{1-\rho^2}\big((1-\rho)\sqrt{1-\lambda^2} + \sqrt{1-\rho^2}\big)\Phi^{-1}(\mu)}{\sqrt{1-\lambda^2}\sqrt{1-\rho^2}\big(\rho + \sqrt{1-\lambda^2}\sqrt{1-\rho^2}\big)}\right) \tag{1.17}$$

$$\text{and} \quad k := \sqrt{1-\rho^2}\big(\rho\sqrt{1-\lambda^2} - \sqrt{1-\rho^2}\big)\,. \tag{1.18}$$

Then $H'(s) = -k\, H_t(s|s)$, where $H_t(t|s) := \partial H(t|s)/\partial t$.

Proof. See proof 1.A.3 in appendix 1.A □

Recall that $G^{-1}(t)$ is such that, *in equilibrium*, a taxpayer who received a signal t will be audited if the tax authority's type is below $G^{-1}(t)$ and let off if it is above $G^{-1}(t)$. Also, $F\big(G^{-1}(t)|s\big)$ is the probability that the tax authority's tolerance is less than or equal to $G^{-1}(t)$ conditional on having received a signal s. Therefore, $H(t|s) := F\big(G^{-1}(t)|s\big)$ is the posterior probability of audit for a taxpayer who receives a signal s, *but who adopts the equilibrium strategy of someone who received a signal* t. Hence, $H(s)$ is the equilibrium posterior probability of audit for a taxpayer who receives a signal s.

We are now ready to characterise the pure strategy equilibria.

Proposition 1.8 Let $\rho^* := \sqrt{2-\lambda^2}$. If $\rho \leq \rho^*$, full compliance is the only pure strategy equilibrium. If $\rho > \rho^*$, for any $\bar{x} \in [0, 1]$, there is a pure strategy equilibrium with

$$x^*(s) = \begin{cases} \bar{x}(1-(1+p)H(s))^{1/k} & \text{if } s > \underline{s} \\ 0 & \text{if } s \leq \underline{s}, \end{cases} \tag{1.19}$$

where $\underline{s} = H^{-1}(1/(1+p))$ and H and k are defined in lemma 1.5.

Proof. Let $\bar{x}$ be the maximum amount of evasion in equilibrium. We know from lemma 1.4 that x^* is increasing for all s such that $x^*(s) > 0$, so that the ordinary inverse $x^{*-1}(x)$ exists for all $x \in (0, \bar{x}]$. Then, given that all taxpayers are following the pure strategy x^*, the payoff of a taxpayer with signal s for $x \in (0, \bar{x}]$ is

$$u_s(x) = x\left(\frac{1}{1+p} - \Pr\left(x \geq \theta(T^s)\right)\right) = x\left(\frac{1}{1+p} - \Pr\left(x \geq x^*(G(T^s))\right)\right) \quad \text{by lemma 1.3} \tag{1.20}$$

$$= x\left(\frac{1}{1+p} - \Pr\left(G^{-1}(x^{*-1}(x)) \geq T^s\right)\right) = x\left(\frac{1}{1+p} - F\left(G^{-1}(x^{*-1}(x))|s\right)\right) \tag{1.21}$$

$$= x\left(\frac{1}{1+p} - H(x^{*-1}(x)|s)\right). \tag{1.22}$$

Assuming that x^* is differentiable, we differentiate eq. (1.22) with respect to x to obtain

$$u_s'(x) = \frac{1}{1+p} - H(x^{*-1}(x)|s) - \frac{x}{x^{*\prime}(x^{*-1}(x))} H_t(x^{*-1}(x)|s). \tag{1.23}$$

Assuming further that $u_s(x)$ has a unique interior maximum, for $x^*(s)$ to be the optimal level of evasion for a taxpayer with signal s, the first-order condition $u_s'(x^*(s)) = 0$ must hold, i.e.,

$$u_s'(x^*(s)) = \frac{1}{1+p} - H(x^{*-1}(x^*(s))|s) - \frac{x^*(s)}{x^{*\prime}(x^{*-1}(x^*(s)))} H_t(x^{*-1}(x^*(s))|s) \tag{1.24}$$

$$= \frac{1}{1+p} - H(s|s) - \frac{x^*(s)}{x^{*\prime}(s)} H_t(s|s) = 0. \tag{1.25}$$

Rearranging eq. (1.25), we have the differential equation

$$\frac{x^{*\prime}(s)}{x^*(s)} = \frac{H_t(s|s)}{1/(1+p) - H(s|s)} = \frac{1}{k}\frac{H'(s)}{1/(1+p) - H(s)} \quad \text{by lemma 1.5} \tag{1.26}$$

$$= \frac{1}{k}\frac{d}{ds}\left[\log(1/(1+p) - H(s))\right]. \tag{1.27}$$

Integrating both sides of eq. (1.27) by s, we have

$$\log x^*(s) = \frac{1}{k}\log\left(1/(1+p) - H(s)\right) + A \tag{1.28}$$

$$\iff x^*(s) = A\left(1/(1+p) - H(s)\right)^{1/k}, \tag{1.29}$$

where A a constant of integration that can be determined by a suitable boundary condition.

The function in eq. (1.29) is qualitatively different depending on the sign of k, which in turn depends on whether ρ is less than, equal to or greater than $\rho^* := \sqrt{2 - \lambda^2}$, so we will analyse these three cases separately.

If $\rho < \rho^*$, then $k < 0$ and $H(s)$ is increasing in s, so that there exists an $\underline{s}$ such that $1/(1+p) = H(\underline{s})$. But then $\left(1/(1+p) - H(s)\right)^{1/k}$ increases without bound as $s \to \underline{s}$ so that eq. (1.29) cannot be an equilibrium. A further reason for why the solution given by eq. (1.29) cannot be an equilibrium if $\rho < \rho^*$ is that it will be decreasing in s for some s, contradicting proposition 1.2. Since by proposition 1.4 full compliance is an equilibrium, it must be the only one.

If $\rho = \rho^*$, then $H(s)$ is constant in s, which means that we are in a pooling equilibrium, but by proposition 1.3, the only possibility is full compliance.

Suppose now that $\rho > \rho^*$. Then $H(s)$ is decreasing in s and $H(1) = 0$, so that the maximum of $x^*(s)$ is $x^*(1) = A(1+p)^{-1/k}$, since x^* is non-decreasing in s. Then, letting $\bar{x} = x^*(1)$, we have $A = \bar{x}(1+p)^{-1/k}$, so that

$$x^*(s) = \bar{x}\big(1 - (1+p)H(s)\big)^{1/k}, \tag{1.30}$$

if $H(s) \leq 1/(1+p)$. Since $x^*(\underline{s}) = 0$, where $\underline{s} = H^{-1}(1/(1+p))$, for x^* to be non-decreasing it must be that $x^*(s) = 0$ for all $s < \underline{s}$.

To check that this is indeed an equilibrium, we find $u_s(x)$ by inverting eq. (1.30) to get

$$x^{*-1}(x) = H^{-1}\left(\frac{1-(x/\bar{x})^k}{1+p}\right) \quad \text{for all } x \in (0, \bar{x}]. \tag{1.31}$$

Since $H(s)$ and $H^{-1}(s)$ are both decreasing in s, if $H(s) \geq 1/(1+p)$ then $x^{*-1}(x) > H^{-1}(1/(1+p)) \geq s$ for all $x \in (0, \bar{x}]$. Then $H(x^{*-1}(x)|s) > H(s|s) = H(s) \geq 1/(1+p)$ since $H(t|s)$ is increasing in t by lemma 1.5, so that $u_s(x) < 0$ for all $x \in (0, \bar{x}]$. Hence $x = 0$ is optimal if $s \leq \underline{s}$.

Similarly, we must verify that eq. (1.30) gives the unique maximum of eq. (1.22) for all $s > \underline{s}$. Lemma 1.A.6 in appendix 1.A demonstrates this by showing that if $\rho > \rho^*$, the payoff given x^* as in eq. (1.30) is strictly quasi-concave in x. □

Note that the maximum amount of evasion $\bar{x}$ in eq. (1.19) can be arbitrarily chosen in $[0, 1]$, thus resulting in multiple equilibria. This multiplicity is due to the fact that for any $\bar{x}$, if everyone else is evading up to $\bar{x}$, evading $x > \bar{x}$ results in a probability of audit of $1 > 1/(1+p)$, yielding a negative payoff, meaning that no taxpayer has an incentive to evade more than $\bar{x}$.

Note also that the equilibrium features a positive mass of taxpayers who are fully compliant. These are precisely the taxpayers who receive signals of low tolerance such that their equilibrium probability of audit is greater than $1/(1+p)$.

1.4 Comparative statics

We now turn to determining the effect of varying the two main parameters of the model, namely the accuracy α of the indicator of evasion and the openness ρ that measures the precision of the signal of the tax authority's tolerance.

The first result establishes that under full-secrecy, evasion is decreasing in accuracy.

Proposition 1.9 For $\rho = 0$, if $\alpha_0 < \alpha_1$, the distribution of evasion when $\alpha = \alpha_0$ first-order stochastically dominates the distribution of evasion when $\alpha = \alpha_1$.

Proof. To prove this we can show that $F_X(x)$ is non-decreasing in α for all x. To do this, note that $\beta = \alpha/(1/(1+p) - (1-\alpha)(1-\mathbb{E}T)) = 1/\big((1/(1+p) - (1-\mathbb{E}T))/\alpha + (1-\mathbb{E}T)\big)$ is increasing in α since by assumption 1.5, $1/(1+p) - (1-\mathbb{E}T) > 0$. Then $\frac{1}{\beta}\big(1 - \frac{1-\beta}{x}\big) = \frac{1}{\beta}(1-1/x) + 1/x$, which is increasing in β and therefore in α since $x \leq 1 \iff 1 - 1/x < 0$.

The fact that F is the cdf of T and therefore is non-decreasing means that F^{-1} is, so that if $F_X(x) > 0$ then $F_X(x)$ is indeed non-decreasing. Furthermore, as $\alpha \to 1$, $\underline{x}(\alpha) \to 0$ monotonically, so that if $F_X(x) = 0$, there exists an α_0 such that $F_X(x) > 0$ if $\alpha > a_0$. Lastly, if $F_X(x) = 1$, then $F_X(x)$ stays constant as α increases. Hence $F_X(x)$ is non-decreasing in x for all x, as required. □

Intuitively, this means that as α increases, for any level of evasion x, the mass of taxpayers evading x or less increases, so that evasion decreases monotonically across taxpayers, so that average evasion too decreases monotonically with α.

The following proposition establishes an analogous result for the opposite case of $\rho = 1$.

Proposition 1.10 If $\rho = 1$, for any t, evasion is non-increasing in α.

PROOF. If $t > \frac{1-1/(1+p)}{1-\alpha} \iff \alpha < \frac{1/(1+p)-(1-t)}{1-(1-t)}$, t taxpayers evade $x^* = 1 - \alpha/\big(1/(1+p) - (1-\alpha)(1-t)\big) = 1 - 1/\big((1/(1+p) - (1-t))/\alpha + (1-t)\big)$, which is decreasing in α, while $1-t$ taxpayers evade 1, which is constant. On the other hand, if $\alpha > \frac{1/(1+p)-(1-t)}{1-(1-t)}$, everybody evades 0, which is constant in α. □

We can now compare the two informational regimes of full secrecy and full openness. Since the information structure of the signals is different, we cannot compare *interim* evasion in the two regimes, but must rather compare *ex-ante* evasion, i.e., before any particular realisation of the tax authority's tolerance. The next result shows when we can make this comparison without making any further assumptions.

Proposition 1.11 If $\alpha > 1/(1+p)$, there is full compliance both when $\rho = 0$ and $\rho = 1$. If $1/(1+p) - c \le \alpha \le 1/(1+p)$, where $c = \big((1-\mathbb{E}T)(1-1/(1+p))\big)/\mathbb{E}T$, ex-ante evasion when $\rho = 0$ is less than when $\rho = 1$. Also, there exists an α_0 such that, ex-ante evasion when $\rho = 0$ is greater than when $\rho = 1$ for all $\alpha \le \alpha_0$.

PROOF. For $\rho = 0$, the equilibrium tends to full compliance as α goes to $\frac{1/(1+p)-(1-\mathbb{E}T)}{1-(1-\mathbb{E}T)} = \frac{1}{1+p} - \frac{(1-\mathbb{E}T)(1-1/(1+p))}{\mathbb{E}T}$. On the other hand, for $\rho = 1$, ex-ante full compliance, i.e., zero evasion for all realisations of T, occurs when $\alpha \ge 1/(1+p) > \frac{1}{1+p} - \frac{(1-\mathbb{E}T)(1-1/(1+p))}{\mathbb{E}T}$. Hence in the given interval of α, full compliance is the equilibrium for $\rho = 0$ but not for $\rho = 1$.

For the second part, note that when $\rho = 0$, $\beta \to 0$ as $\alpha \to 0$, so that $F_X(x) \to 0$ for all $x < 1$, thus approaching the full evasion equilibrium. On the other hand, when $\rho = 1$, as $\alpha \to 0$, the equilibrium becomes the classical one where there is full compliance if $1-t > 1/(1+p)$ and full evasion if $1-t < 1/(1+p)$, so that in ex-ante terms, there is less evasion when $\rho = 1$ than when $\rho = 0$. Furthermore, by proposition 1.9 and proposition 1.10, these extremes are reached monotonically, so that there exists some threshold for α below which $\rho = 1$ yields less average evasion than $\rho = 0$. □

This result shows that when comparing the extremes of the informativeness of the taxpayers' signals of tolerance, full secrecy does indeed lead to greater (in fact, full) compliance than complete openness if the tax authority can distinguish taxpayers by evasion with sufficient accuracy. On the other hand, if the accuracy of the indicator is sufficiently low then full openness leads to more compliance in ex-ante terms. The reason for the latter is that if indicator accuracy is too low, the tax authority is essentially auditing at random. In the absence of any information about the current realisation of the tax authority's tolerance, taxpayers must rely solely on their prior beliefs, which, by assumption 1.5 lead them to always evade fully. On the other hand, if they are perfectly informed about the tax authority type, there will be compliance whenever the tax authority's type is low enough to deter evasion.

A similar intuition is behind the next result, which describes how the equilibrium changes as ρ varies continuously when auditing is completely random.

Proposition 1.12 Under assumption 1.6, if $\alpha = 0$, ex-ante evasion is decreasing in ρ.

PROOF. If $\alpha = 0$, for a given t, those with s such that $\mathbb{E}T^s > 1/(1+p) \iff \mathbb{E}T^s < p/(1+p)$ evade 0 and the rest evade 1. Therefore, the mass of compliant taxpayers is $\Pr\left(\mathbb{E}T^S \leq p/(1+p)\right)$ and lemma 1.A.5 shows that this is increasing in ρ if assumption 1.5 holds, giving the result. □

This result shows that, under assumption 1.6, the above logic holds also for intermediate values of ρ: as the signals become more accurate, receiving a signal of low tolerance is more and more informative of a high exogenous probability of audit, leading to more taxpayers complying[23].

At the other opposite extreme, the next result shows how the equilibrium varies as ρ varies continuously when the tax authority has perfect pre-audit information about taxpayers' evasion.

Proposition 1.13 If $\alpha = 1$, for a fixed signal s, equilibrium evasion tends to zero monotonically as $\rho \to \rho^*$ and as $\rho \to 1$.

PROOF. See proof 1.A.4 in appendix 1.A. □

Proposition 1.13 shows that for any given signal, if ρ is sufficiently high, evasion is decreasing in openness. This result is in direct opposition to the intuitive arguments put forward by the IRS as quoted in the introduction, so the intuition behind it merits some discussion. We begin by interpreting the function $k(\rho)$ in the exponent of eq. (1.19). To do this, we use lemma 1.5 to write k as

$$k = \frac{-H'(s)}{H_t(s|s)} = \frac{-H_s(s|s)}{H_t(s|s)} - 1\,, \tag{1.32}$$

where subscripts denote partial derivatives. Recall from eq. (1.22) that $H(t|s)$ is the audit probability for a taxpayer who received a signal s and evades like someone who received a signal t. Therefore, H_t is the rate at which the audit probability rises due to increases in evasion, capturing the incentive to reduce evasion in order to avoid being audited. On the other

[23] Recall that the comparative statics are carried out in ex-ante terms, so ρ does not affect the distribution of signals, which is always given by the prior distribution of S.

Figure 1.3: Plots of equilibrium evasion $x^*(s)$

$s = 0.5, 0.6, 0.7,\ p = 0.2$. Higher curves correspond to higher s.

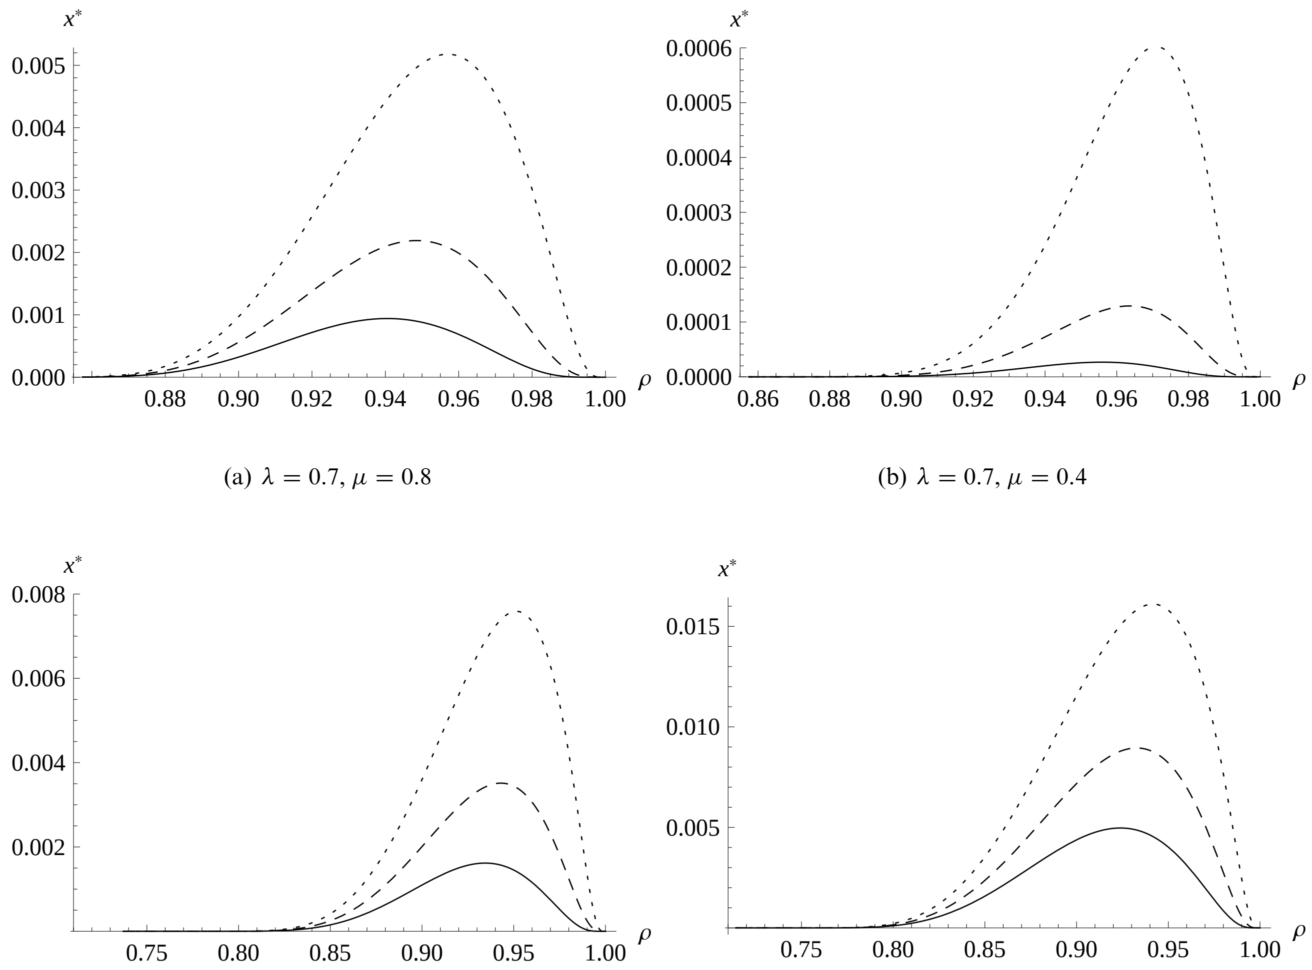

(a) $\lambda = 0.7$, $\mu = 0.8$

(b) $\lambda = 0.7$, $\mu = 0.4$

(c) $\lambda = 0.4$, $\mu = 0.4$

(d) $\lambda = 0$

hand, the term $-H_s(s|s)$ captures a kind of informational rent from the signal. To see this, let $u^*(s) := u_s(x^*(s)) = x^*(s)\big(1/(1+p) - H(s)\big)$ be the equilibrium payoff for a taxpayer who receives a signal s. Then $u^{*\prime}(s) = x^{*\prime}(s)\, u_s'(x^*(s)) - x^*(s) H_s(s|s) = -x^*(s) H_s(s|s)$, where the last equality follows from the fact that $x^*(s)$ satisfies the first-order condition (1.23). Therefore, the greater the value of $-H_s(s|s)$, the greater the benefit of having a higher signal.

Having expressed k as a function of the ratio of $-H_s$ and H_t, we need to understand how changes in ρ affect these two. Firstly, under assumption 1.6, $-H_s(s|s)$ and $H_t(s|s)$ are $\rho\, h(s,\rho)$ and $1/\big(\sqrt{1-\lambda^2}\sqrt{1-\rho^2}+\rho\big)\, h(s,\rho)$, respectively, where h is a common term that we will ignore. On the one hand, we see that $-H_s$ is increasing in ρ, as we would expect; as the precision of the signals increases, their informational content increases, so that taxpayers can gain more from receiving a signal of high tolerance. Since this is a zero-sum game, this gain must come at the expense of the tax authority, so that we can interpret the fact that $-H_s$ is increasing in ρ as embodying the fear of openness exemplified in the passage quoted in the introduction.

On the other hand, H_t, the incentive to reduce evasion, also increases in ρ, for $\rho > \rho^*$; as the spread of the signals decreases, taxpayers become more homogeneous, since the signals were the only source of heterogeneity. Recall from proposition 1.3 that when taxpayers pool, it is optimal for them to slightly undercut each other in order to escape audit. As taxpayers become more homogeneous, this Bertrand-like competition intensifies, reducing the incentives to evade. Thus, increasing openness ρ has a beneficial effect for the tax authority that arises due to the strategic interaction between taxpayers and which counteracts the more obvious detrimental one described earlier.

The ratio of $-H_s$ and H_t measures the relative strength of these two effects. As ρ increases, $-H_s$ increases linearly, whereas H_t increases at an increasing rate[24], so that their ratio, and therefore $k(\rho)$, first increases and then decreases with ρ. Since the reciprocal of k appears as the exponent of a value less than 1 in eq. (1.19), x^* would increase and then decrease in ρ, *if $H(s)$ were kept constant*. The only difficulty here is that $H(s)$ depends on ρ also, and it is possible for it to decrease when $1/k(\rho)$ decreases, so that it is not analytically clear that x^* first increases and then decreases in ρ as it would if the only dependence on ρ were through $k(\rho)$. Nonetheless, the proof of proposition 1.13 shows that whenever the effects of ρ on x^* through H and k are opposed, the effect through k dominates, so that the intuition outlined above is indeed the driving force of the result. In fact, numerical simulations[25] suggest that x^* is indeed unimodal for a wide range of parameters, even though the exact conditions are at present undetermined. The examples in fig. 1.3 illustrate this point.

1.5 Discussion

So far, we have characterised the equilibria and analysed their comparative statics along the edges of the unit square formed by the parameters α and ρ, but what can we say about intermediate values of both α and ρ? A natural starting point is to extend the approach used in the construction of the equilibrium in the case of a perfect indicator in section 1.3.3 by allowing α to take intermediate values. Unfortunately, we cannot continue assuming that taxpayers adopt pure strategies because the solution to the suitably generalised FOC in eq. (1.23) does not constitute an equilibrium. Figure 1.4 shows an example of this: the solid line depicts u^*, the payoff obtained by following an evasion strategy given by the FOC, whereas the dashed line shows the payoff obtained by evading $x = 1$ when everyone else is following the FOC strategy; for some values of the signal s, deviating to evading 1 yields a strictly higher payoff than the one obtained by following the FOC strategy.

Given that the equilibrium for $\rho = 1$ and low values of α features mixed strategies, it seems likely that the equilibrium in the interior of the ρ-α space will also feature mixed strategies, but the solution to this general problem remains elusive.

Nonetheless, the comparative statics result for the perfect indicator case (proposition 1.13) shows that the competitive argument behind proposition 1.3 that makes full compliance the

[24]This is because $\partial[\sqrt{1-\lambda^2}\sqrt{1-\rho^2}+\rho]/\partial\rho = 1 - \sqrt{1-\lambda^2}\rho/\sqrt{1-\rho^2}$ is decreasing in ρ, and the reciprocal of a concave function is convex.

[25]An interactive Mathematica notebook for exploring the solution x^* numerically for any parameter value is available from the author.

Figure 1.4: Insufficiency of first order condition when $\alpha \neq 1$

Solid line is u^* according to the FOC and dashed line is $u(1)$. $\rho = 0.8$, $\alpha = 0.5$ and $p = 0.2$

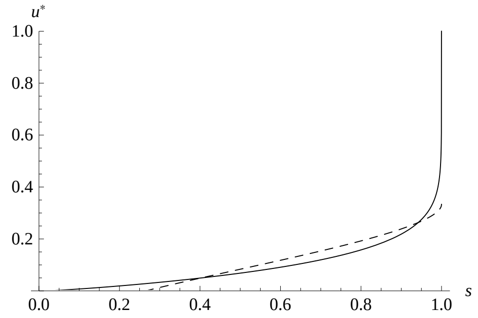

only possible pooling equilibrium exerts its force even when taxpayers are informationally heterogeneous. This Bertrand-like mechanism is novel in the literature on tax evasion and is unlikely to have no relevance when the tax authority's ability to select the most egregious evaders for audit is less than perfect.

A criticism that could be levelled at the equilibrium in the perfect indicator case is that it is not unique, unlike the ones usually found in the literature on tax evasion. This multiplicity is natural in coordination games, although not so in global games which typically feature a unique equilibrium, but one can argue that the multiple equilibria are *qualitatively* similar: in particular, the relative level of evasion across taxpayers is the same in all of them, since the multiplicity derives from the choice of a multiplicative normalisation constant. Indeed, all equilibria feature exactly the same comparative statics and therefore yield the same conclusions about the desirability or otherwise of secrecy.

We saw earlier that the fact that individuals are homogeneous except for receiving private signals about a fundamental type makes the model considered in this paper a global game. The contribution to the literature in this area is that this paper considers a structure for the signal that is radically different from the additive error one commonly used. This comes at the cost of some analytical tractability, but modelling the joint distribution of the type and the signal makes it possible to consider the full range of degrees of stochastic dependence between the two, from full dependence or comonotonicity to independence, something which cannot be done with an additive error.

The other innovation of this model is the introduction of an indicator of evasion that the tax authority can use to select taxpayers for audit. As mentioned in the introduction, this represents a departure from existing models where such information is either absent or about incomes rather than directly about evasion. One important assumption that the present paper makes is that the process by which the indicator is generated is known. In reality, taxpayers are not very well informed about *how*[26] as well as *how many* individuals are selected for audit.

[26] Indeed the IRS goes to great lengths to hide details of how DIF is computed, as the following quote from FTC (2011) shows:

> The IRS statistical technique known as discriminant function scores (DIF) is protected from disclosure under the investigatory records exception discussed above. DIF scores are used by IRS to calculate tolerances outside of which audits become more likely. They also are used to determine whether a fraud investigation should be undertaken and to identify collection methods. Discrimination function (DIF) scores are exempt from FOIA because they are investigative techniques, the

This could mean that there is uncertainty about α as well as the t considered here.

Perhaps a more fundamental objection to the model presented here is that the taxpayers are assumed to be homogeneous except for their signals, so that they all owe the same amount of taxes. This assumption is obviously not realistic, even within a fairly narrow class of taxpayers, but it is a useful one, in that it highlights how heterogeneity in information drives evasion behaviour, rather than the extensively studied heterogeneity in type studied by, for example, Reinganum and Wilde (1986) and Sanchez and Sobel (1993).

Lastly, how applicable are the results of the current model to contexts other than tax evasion? In principle, the model could be applied to a general regulatory context similar to the one in Hansen, Krarup, and Russell (2006). The limiting factors are the features specific to the taxation context, namely the presence of a very large number of agents to be audited and the very specific proportional penalty structure. The first feature is integral to the model because it ensures that an individual cannot alone influence the aggregate, which might not be the case in scenarios where there is a small number of "big" players, such as large firms. As for the second feature, the penalty structure could easily be modified to fit a different context, say to accommodate a fixed, rather than proportional fine. We can see this by noting that the results, especially proposition 1.13 carry through even if we let $p = 0$.

1.6 Conclusion

This paper has demonstrated in proposition 1.11 that the relative desirability of policies of extreme secrecy or openness about how tolerant a tax authority is towards evasion depends crucially on the accuracy of the information it uses to selects individuals for audit. Although each type of information has been modelled separately in the literature on tax evasion, this is the first paper to allow both, generating such a result.

Furthermore, proposition 1.3 and proposition 1.13 show how Bertrand-like competition among taxpayers to reduce their evasion in order to reduce the probability with which they are audited improves compliance, with greater information about the tax authority's tolerance leading to greater compliance. This competitive mechanism is novel and provides a counter-argument to the view that a tax agency must strive for secrecy about itself in order to foster compliance.

This paper could usefully be extended in three directions. Firstly, one could relax the assumption that the technology used by the tax authority to assess whom to audit is known to the taxpayers. Secondly, the homogeneity of taxpayers in every respect other than their information about the tax authority could be relaxed to allow evasion differ by factors other than information. Thirdly, alternative penalty structures, such as fixed or affine ones could be considered, in order to improve direct applicability of the results to contexts other than taxation, where proportional fines are the norm.

release of which would harm IRS ability to enforce the tax laws. E.g., such release could enable taxpayers to develop plans to avoid flagging of their returns for audit and thus circumvent law enforcement.

Appendix 1.A Proofs and mathematical results

Proof 1.A.1 (Proof of lemma 1.1) Firstly, note that because the conditional distribution of Y given $X = x$ is discontinuous at $y = x$, we will use the Stieltjes integral[27] throughout this proof.

First we find the joint distribution of X and Y in order to show its symmetry. First we rewrite eq. (1.2) as

$$\Pr(Y \leq y \mid X = x) = \begin{cases} (1-\alpha)F_X(y) & \text{if } x > y \\ (1-\alpha)F_X(y) + \alpha & \text{if } x \leq y \end{cases} \tag{1.33}$$

$$= (1-\alpha)F_X(y) + \alpha \mathbb{1}_{[0,y]}(x)\,, \tag{1.34}$$

where $\mathbb{1}_A(x)$ equals 1 if $x \in A$ and 0 otherwise. Then the joint probability that $Y \leq y$ and $X \leq x$ is

$$\Pr(Y \leq y, X \leq x) = \int_0^x \Pr(Y \leq y \mid X = z)\, dF_X(z) \tag{1.35}$$

$$= (1-\alpha)\int_0^x F_X(y)\, dF_X(z) + \alpha \int_0^x \mathbb{1}_{[0,y]}(z)\, dF_X(z) \tag{1.36}$$

$$= (1-\alpha)F_X(y)F_X(x) + \alpha \int_0^{\min(x,y)} F_X(z)\, d \tag{1.37}$$

$$= (1-\alpha)F_X(y)F_X(x) + \alpha F_X(\min(x,y)) \tag{1.38}$$

which is symmetric in x and y.

To find $\mathrm{Corr}(X,Y) = (\mathbb{E}(XY) - \mathbb{E}(X)\mathbb{E}(Y))/\sqrt{\mathrm{Var}(X)\,\mathrm{Var}(Y)}$, we first use the law of iterated expectations to write $\mathbb{E}(XY) = \mathbb{E}(X\mathbb{E}(Y|X))$. Next, rewrite $\Pr(Y \leq y \mid X = x)$ as a function of y and write $F_x(y) := (1-\alpha)F_X(y) + \alpha\mathbb{1}_{[x,1]}(y)$, so that

$$\mathbb{E}(Y \mid X = x) = \int_0^1 y\, dF_x(y) = (1-\alpha)\int_0^1 y\, dF_X(y) + \alpha \int_0^1 y\, d\mathbb{1}_x(y) \tag{1.39}$$

$$= (1-\alpha)\mathbb{E}X + \alpha x\,, \tag{1.40}$$

so that $\mathbb{E}(XY) = \mathbb{E}[X\big((1-\alpha)\mathbb{E}(X) + \alpha X\big)] = (1-\alpha)\mathbb{E}(X)^2 + \alpha\mathbb{E}(X^2)$, and hence $\mathbb{E}(XY) - \mathbb{E}(X)\mathbb{E}(Y) = \alpha(\mathbb{E}(X^2) - \mathbb{E}(X)^2) = \alpha \mathrm{Var}\, X$. Then $\mathrm{Corr}(X,Y) = \mathrm{Cov}(X,Y)/\sqrt{\mathrm{Var}\, X\, \mathrm{Var}\, Y} = \mathrm{Cov}(X,Y)/\mathrm{Var}\, X = \alpha$, as required. ■

Lemma 1.A.1 Under assumption 1.6, the cdf of $T|S = s$ is

$$F(t|s) = \Phi\left(\frac{\Phi^{-1}(t) - \rho\Phi^{-1}(s) - \lambda(1-\rho)\Phi^{-1}(\mu)}{\sqrt{1-\lambda^2}\sqrt{1-\rho^2}}\right), \tag{1.41}$$

and the marginal cdf of S and T is

$$F(s) = \Phi\left(\frac{\Phi^{-1}(s) - \lambda\Phi^{-1}(\mu)}{\sqrt{1-\lambda^2}}\right). \tag{1.42}$$

[27] See Ok (2011) for details.

PROOF. Recall that if $\tilde{S}$ and $\tilde{T}$ are standard bivariate normal with correlation ρ, then $\tilde{T}|\tilde{S} = s$ is normal with mean ρ^* and variance $1-\rho^2$, so that $\Pr(\tilde{T} \le t|\tilde{S} = s) = \Phi\big((t-\rho^*)/\sqrt{1-\rho^2}\big)$. Then,

$$\Pr(T \le t|S = s) = \Pr(\Phi(\sqrt{1-\lambda^2}\tilde{T} + \lambda\Phi^{-1}(\mu)) \le t|\Phi(\sqrt{1-\lambda^2}\tilde{S} + \lambda\Phi^{-1}(\mu)) = s) \tag{1.43}$$

$$= \Pr\left(\tilde{T} \le \frac{\Phi^{-1}(t) - \lambda\Phi^{-1}(\mu)}{\sqrt{1-\lambda^2}}\middle|\tilde{S} \le \frac{\Phi^{-1}(s) - \lambda\Phi^{-1}(\mu)}{\sqrt{1-\lambda^2}}\right) \tag{1.44}$$

$$= \Phi\left(\frac{\Phi^{-1}(t) - \lambda\Phi^{-1}(\mu) - \rho(\Phi^{-1}(s) - \lambda\Phi^{-1}(\mu))}{\sqrt{1-\lambda^2}\sqrt{1-\rho^2}}\right) \tag{1.45}$$

$$= \Phi\left(\frac{\Phi^{-1}(t) - \rho\Phi^{-1}(s) - \lambda(1-\rho)\Phi 1(\mu)}{\sqrt{1-\lambda^2}\sqrt{1-\rho^2}}\right), \tag{1.46}$$

as required. The marginal cdf is then obtained by setting $\rho = 0$ above. □

Lemma 1.A.2 Under assumption 1.6, $F(t|s)$ satisfies assumption 1.4.

PROOF. Since $F(t|s)$ has a density $f(t|s)$, we will show this by proving that $f(t|s)$ is log-supermodular in s, t, which in turn is equivalent to showing that $\partial^2 \log f(t|s)/\partial t \partial s > 0$. Firstly, the density is

$$f(t|s) = \frac{\partial}{\partial t}F(t|s) = \frac{1}{\sqrt{1-\lambda^2}\sqrt{1-\rho^2}} \times \phi\left(\frac{\Phi^{-1}(t) - \rho\Phi^{-1}(s) - \lambda(1-\rho)\Phi^{-1}(\mu)}{\sqrt{1-\lambda^2}\sqrt{1-\rho^2}}\right)/\phi(\Phi^{-1}(t)) \tag{1.47}$$

$$\implies \log f(t|s) = \log\phi\left(\frac{\Phi^{-1}(t) - \rho\Phi^{-1}(s) - \lambda(1-\rho)\Phi^{-1}(\mu)}{\sqrt{1-\lambda^2}\sqrt{1-\rho^2}}\right) + C, \tag{1.48}$$

where C is a function of t alone. Since $\phi(x) = \exp(-x^2/2)$,

$$\frac{\partial^2}{\partial t \partial s}\log f(t|s) = \frac{\partial^2}{\partial t \partial s}\left[-\frac{\big(\Phi^{-1}(t) - \rho\Phi^{-1}(s) - \lambda(1-\rho)\Phi^{-1}(\mu)\big)^2}{2\big(\sqrt{1-\lambda^2}\sqrt{1-\rho^2}\big)^2}\right] \tag{1.49}$$

$$= \frac{\partial}{\partial t}\left[\rho\frac{\big(\Phi^{-1}(t) - \rho\Phi^{-1}(s) - \lambda(1-\rho)\Phi^{-1}(\mu)\big)}{\phi(\Phi^{-1}(s))\big(\sqrt{1-\lambda^2}\sqrt{1-\rho^2}\big)^2}\right] \tag{1.50}$$

$$= \frac{\rho}{\phi(\Phi^{-1}(t))\phi(\Phi^{-1}(s))\big(\sqrt{1-\lambda^2}\sqrt{1-\rho^2}\big)^2} > 0, \tag{1.51}$$

as required. □

Lemma 1.A.3 If $\Pr(X \le x) = \Phi\big((\Phi^{-1}(x) - a)/b\big)$, then $\mathbb{E}X = \Phi\big(a/\sqrt{1+b^2}\big)$.

PROOF. If $\Pr(X \le x) = \Phi\big((\Phi^{-1}(x) - a)/b\big)$, then $X = \Phi(a + bY)$, where Y is standard normal, since $\Pr(X \le x) = \Phi\big((\Phi^{-1}(x) - a)/b\big) = \Pr\big(Y \le (\Phi^{-1}(x) - a)/b\big) = \Pr(\Phi(a + bY) \le x)$.

Then, $\mathbb{E}X = \mathbb{E}[\Phi(a+bY)] = \mathbb{E}[\Pr(Z \le a+bY \mid Y)] = \Pr(Z \le a+bY) = \Pr\big((Z - bY)/\sqrt{1+b^2} \le a/\sqrt{1+b^2}\big) = \Phi\big(a/\sqrt{1+b^2}\big)$, where Z is standard normal. □

Lemma 1.A.4 Under assumption 1.6, $\mathbb{E}T = \Phi\big(\lambda\Phi^{-1}(\mu)/\sqrt{2-\lambda^2}\big)$.

PROOF. Follows from lemma 1.A.1 and lemma 1.A.3 □

Lemma 1.A.5 Under assumption 1.6, $\Pr\big(\mathbb{E}T^S < p/(1+p)\big)$ is increasing in ρ.

PROOF. Combining lemma 1.A.1 and lemma 1.A.3, we have

$$\mathbb{E}T^S = \mathbb{E}(T|S) = \Phi\left(\frac{\rho\Phi^{-1}(S) + \lambda(1-\rho)\Phi^{-1}(\mu)}{\sqrt{1+(1-\lambda^2)(1-\rho^2)}}\right), \tag{1.52}$$

so that $\Pr\big(\mathbb{E}T^S < p/(1+p)\big)$ is

$$= \Pr\left(S \le \Phi\big(\sqrt{1+(1-\lambda^2)(1-\rho^2)}\Phi^{-1}(p/(1+p)) - (1-\rho)\lambda\Phi^{-1}(\mu)\big)/\rho\right) \tag{1.53}$$

$$= \Phi\left(\frac{\sqrt{1+(1-\lambda^2)(1-\rho^2)}\Phi^{-1}(p/(1+p)) - \lambda\Phi^{-1}(\mu)}{\rho\sqrt{1-\lambda^2}}\right) \tag{1.54}$$

$$= \Phi\left(\sqrt{1+\frac{\lambda^2}{1-\lambda^2}\frac{1}{\rho^2}}\Phi^{-1}(p/(1+p)) - \frac{\lambda}{\sqrt{1-\lambda^2}}\frac{1}{\rho}\Phi^{-1}(\mu)\right) \tag{1.55}$$

$$= \Phi\left(\sqrt{1+\ell(\rho)^2}\Phi^{-1}(p/(1+p)) - \ell(\rho)\Phi^{-1}(\mu)\right), \tag{1.56}$$

where $\ell(\rho) = \lambda/\sqrt{1-\lambda^2}\,1/\rho$. We need to show that the term inside Φ in eq. (1.56) is increasing in ρ and we do this by finding its derivative, which is

$$\ell'(\rho)\left(\frac{1}{\sqrt{1+\ell(\rho)^2}}\Phi^{-1}(p/(1+p)) - \Phi^{-1}(\mu)\right), \tag{1.57}$$

which is increasing in ρ iff

$$\Phi^{-1}(p/(1+p)) < \sqrt{1+\ell(\rho)^2}\Phi^{-1}(\mu), \tag{1.58}$$

since $\ell(\rho)$ is decreasing in ρ. Under assumption 1.5, $1 - \mathbb{E}T < 1/(1+p)$ which implies

$$\Phi(p/(1+p)) < \frac{\lambda}{\sqrt{2-\lambda^2}}\Phi^{-1}(\mu) \quad \text{by lemma 1.A.4} \tag{1.59}$$

$$< \Phi^{-1}(\mu) < \sqrt{1+\ell(\rho)^2}\Phi^{-1}(\mu), \tag{1.60}$$

as required. □

Proof 1.A.2 (Proof of lemma 1.2) We invert $F(t|s)$ from lemma 1.A.1, to obtain

$$F^{-1}(t|s) = \Phi\left(\sqrt{1-\lambda^2}\sqrt{1-\rho^2}\Phi^{-1}(t) + \rho\Phi^{-1}(s) + \lambda(1-\rho)\Phi^{-1}(\mu)\right) \tag{1.61}$$

$$\implies G(t) := F^{-1}(t|t) = \Phi\left(\big(\sqrt{1-\lambda^2}\sqrt{1-\rho^2} + \rho\big)\Phi^{-1}(t) + \lambda(1-\rho)\Phi^{-1}(\mu)\right), \tag{1.62}$$

which is strictly increasing and has an inverse

$$G^{-1}(t) = \Phi\left(\frac{\Phi^{-1}(t) - \lambda(1-\rho)\Phi^{-1}(\mu)}{\sqrt{1-\lambda^2}\sqrt{1-\rho^2}+\rho}\right) \tag{1.63}$$

for any $t \in [0,1]$ if $\lambda < 1$. ■

Proof 1.A.3 (Proof of lemma 1.5) First we compute $H(t|s) := F(G^{-1}(t)|s) = \Phi(\square)$ and

$$H_t(t|s) := \frac{\partial H(t|s)}{\partial t} = \frac{1}{\sqrt{1-\lambda^2}\sqrt{1-\rho^2}\big(\sqrt{1-\lambda^2}\sqrt{1-\rho^2}+\rho\big)} \frac{\phi(\square)}{\phi(\Phi^{-1}(t))} \tag{1.64}$$

$$\text{where} \quad \square := \frac{\begin{pmatrix} \Phi^{-1}(t) - \rho\big(\sqrt{1-\lambda^2}\sqrt{1-\rho^2}+\rho\big)\Phi^{-1}(s) \\ -\lambda\sqrt{1-\rho^2}\big(\sqrt{1-\rho^2}+(1-\rho)\sqrt{1-\lambda^2}\big)\Phi^{-1}(\mu) \end{pmatrix}}{\sqrt{1-\lambda^2}\sqrt{1-\rho^2}\big(\sqrt{1-\lambda^2}\sqrt{1-\rho^2}+\rho\big)}. \tag{1.65}$$

Next, we compute $H(s) := H(s|s) = \Phi(\blacksquare)$ and

$$H'(s) = \frac{-k}{\sqrt{1-\lambda^2}\sqrt{1-\rho^2}\big(\sqrt{1-\lambda^2}\sqrt{1-\rho^2}+\rho\big)} \frac{\phi(\blacksquare)}{\phi(\Phi^{-1}(s))} \tag{1.66}$$

$$\text{where} \quad \blacksquare := \frac{-k\Phi^{-1}(s) - \lambda\sqrt{1-\rho^2}\big(\sqrt{1-\rho^2}+(1-\rho)\sqrt{1-\lambda^2}\big)\Phi^{-1}(\mu)}{\sqrt{1-\lambda^2}\sqrt{1-\rho^2}\big(\sqrt{1-\lambda^2}\sqrt{1-\rho^2}+\rho\big)} \tag{1.67}$$

$$\text{and} \quad k := \sqrt{1-\rho^2}\big(\rho\sqrt{1-\lambda^2} - \sqrt{1-\rho^2}\big). \tag{1.68}$$

If $t = s$, then $\blacksquare = \square$, so that $H'(s) = -k\, H_t(s|s)$, as required. ■

Lemma 1.A.6 If $\rho > 1/\sqrt{2-\lambda^2}$, given x^* in eq. (1.19), $u_s(x)$ is strictly quasi-concave in x.

Proof. Firstly we need to invert $x^*(s)$ in eq. (1.19) to obtain

$$x^{*-1}(x) = H^{-1}\left(\frac{1-x^k}{1+p}\right). \tag{1.69}$$

Furthermore,

$$H(H^{-1}(t)|s) = \Phi\left[-\frac{1}{k}\Phi^{-1}(t) - \frac{\rho}{\sqrt{1-\lambda^2}\sqrt{1-\rho^2}}\Phi^{-1}(s) + C\Phi^{-1}(\mu)\right], \tag{1.70}$$

$$\text{where} \quad C := \frac{\lambda\rho\big(1+\rho+\sqrt{1-\lambda^2}\sqrt{1-\rho^2}\big)}{(1+\rho)\sqrt{1-\lambda^2}\big(\sqrt{1-\rho^2}-\rho\sqrt{1-\lambda^2}\big)}, \tag{1.71}$$

so that,

$$u_s(x) = x\left(\frac{1}{1+p} - H(x^{*-1}(x)|s)\right) \tag{1.72}$$

$$= x\left(\frac{1}{1+p} - \Phi\left(-\frac{1}{k}\Phi^{-1}\left(\frac{1-x^k}{1+p}\right) + D\right)\right), \tag{1.73}$$

$$\text{where} \quad D := -\frac{\rho}{\sqrt{1-\lambda^2}\sqrt{1-\rho^2}}\Phi^{-1}(s) + C\Phi^{-1}(\mu). \tag{1.74}$$

Then,

$$\log u_s(x) = \log x + \log\left(\frac{1}{1+p} - \Phi\left(-\frac{1}{k}\Phi^{-1}\left(\frac{1-x^k}{1+p}\right) + D\right)\right) \tag{1.75}$$

$$\implies \frac{d}{dx}\log u_s(x) = \frac{1}{x} - \frac{x^{k-1}}{1+p}\frac{\phi\left(-\frac{1}{k}\Phi^{-1}\left(\frac{1-x^k}{1+p}\right)+D\right)/\phi\left(\Phi^{-1}\left(\frac{1-x^k}{1+p}\right)\right)}{\frac{1}{1+p} - \Phi\left(-\frac{1}{k}\Phi^{-1}\left(\frac{1-x^k}{1+p}\right)+D\right)} \tag{1.76}$$

$$= \frac{1}{x}\left(1 - \frac{\phi(-z/k+D)}{\frac{1}{1+p} - \Phi(-z/k+D)}\Bigg/\frac{\phi(z)}{\frac{1}{1+p} - \Phi(z)}\right), \tag{1.77}$$

$$\text{where} \quad z = \Phi^{-1}\left(\frac{1-x^k}{1+p}\right). \tag{1.78}$$

Note that $\phi(x)/(1/(1+p) - \Phi(x)) = (1-\Phi(x))/(1/(1+p) - \Phi(x))\phi(x)/(1-\Phi(x))$ is increasing since the hazard rate $\phi(x)/(1-\Phi(x))$ of the standard normal distribution is increasing and so is $(1-\Phi(x))/(1/(1+p) - \Phi(x))$ since $1/(1+p) < 1$. Also, since $k > 0$ if $\rho > \rho^*$, z is decreasing in x, $\phi(-z/k+D)/(1/(1+p) - \Phi(-z/k+D))$ is increasing and $\phi(z)/(1/(1+p) - \Phi(z))$ is decreasing in x, so that $d\log u_s(x)/dx$ is decreasing in x. Hence $\log u_s(x)$ is strictly concave and therefore strictly quasi-concave. Since strict quasi-concavity is preserved by increasing transformations, it follows that $u_s(x)$ is also strictly quasi-concave.□

Proof 1.A.4 (Proof of proposition 1.13) Since we will be interested in the behaviour of x^* as a function of ρ, let us rewrite the equilibrium amount of evasion as

$$\xi(\rho) := \begin{cases} \psi(\rho)^{1/k(\rho)} & \text{if } \psi(\rho) > 0 \\ 0 & \text{if } \psi(\rho) \leq 0, \end{cases} \tag{1.79}$$

$$\text{where} \quad \psi(\rho) := \Phi\big(\eta(\rho)\big) - p/(1+p), \tag{1.80}$$

$$\eta(\rho) := a(\rho)\Phi^{-1}(s) + b(\rho)\lambda\Phi^{-1}(\mu) \tag{1.81}$$

$$= a(\rho)\big(\Phi^{-1}(s) - \lambda\Phi^{-1}(\mu)\big) + \frac{\lambda\Phi^{-1}(\mu)}{\rho + \sqrt{1-\lambda^2}\sqrt{1-\rho^2}}, \tag{1.82}$$

$$\text{and} \quad a(\rho) := \frac{1}{\sqrt{1-\lambda^2}}\frac{\sqrt{1-\lambda^2}\rho - \sqrt{1-\rho^2}}{\rho + \sqrt{1-\lambda^2}\sqrt{1-\rho^2}}, \tag{1.83}$$

$$b(\rho) := \frac{1}{\sqrt{1-\lambda^2}}\frac{\sqrt{1-\lambda^2}(1-\rho) + \sqrt{1-\rho^2}}{\rho + \sqrt{1-\lambda^2}\sqrt{1-\rho^2}} \tag{1.84}$$

$$= \frac{1}{\rho + \sqrt{1-\lambda^2}\sqrt{1-\rho^2}} - a(\rho). \tag{1.85}$$

Note that the expression $\xi(\rho)$ differs from eq. (1.19) in proposition 1.8 only in that the term ψ defined in eq. (1.80) is of the form in eq. (1.29), without the normalisation $A = \bar{x}(1+p)^{-1/k}$, which is not required here since we are concerned with the comparative statics with respect to ρ and don't require the maximum amount of evasion to be a particular fixed value.[28]

[28] In any case, $(1+p)^{-1/k}$ is increasing in ρ for low ρ and decreasing in ρ for high ρ, reinforcing the result.

To prove the result we first note that since $k(\rho) \to 0$ as $\rho \to \rho^*$ or $\rho \to 1$, and $\psi(\rho) < 1$, it follows that $\xi(\rho) \to 0$ as $\rho \to \rho^*$ or $\rho \to 1$. It remains to be proven that this convergence is monotone.

We will find it convenient to work with $\tilde{\xi}$, defined by $\tilde{\xi}(\rho) := \log k(\rho) + \tilde{\psi}(\rho)$, where $\tilde{\psi}(\rho) := -\log\big(-\log\psi(\rho)\big)$. Since $-\log(-\log(x))$ is increasing in x, showing that ξ is decreasing is equivalent to showing that $\tilde{\xi}$ is decreasing, so we compute $\tilde{\xi}'(\rho) = (\log k)'(\rho) + \tilde{\psi}'(\rho)$, where

$$(\log k)'(\rho) = -\frac{\rho}{1-\rho^2} + \frac{C_1(\rho)}{\sqrt{1-\rho^2}} \tag{1.86}$$

$$\text{and} \quad C_1(\rho) := \frac{\rho + \sqrt{1-\lambda^2}\sqrt{1-\rho^2}}{\sqrt{1-\lambda^2}\rho - \sqrt{1-\rho^2}}, \tag{1.87}$$

$$\text{and} \quad \tilde{\psi}'(\rho) = \frac{\psi'(\rho)}{-\psi(\rho)\log\psi(\rho)} = \frac{\phi\big(\eta(\rho)\big)}{-\psi(\rho)\log\psi(\rho)}\frac{C_2(\rho)}{\sqrt{1-\rho^2}} \tag{1.88}$$

$$\text{where} \quad C_2(\rho) := \frac{(2-\lambda^2)\big(\Phi^{-1}(s) - \lambda\Phi^{-1}(\mu)\big) + \big(\sqrt{1-\lambda^2}\rho - \sqrt{1-\rho^2}\big)\lambda\Phi^{-1}(\mu)}{\sqrt{1-\lambda^2}\big(\rho + \sqrt{1-\lambda^2}\sqrt{1-\rho^2}\big)^2}. \tag{1.89}$$

$$\text{so that} \quad \tilde{\xi}'(\rho) = -\frac{\rho}{1-\rho^2} + \frac{C_1(\rho)}{\sqrt{1-\rho^2}} + \frac{\phi\big(\eta(\rho)\big)}{-\psi(\rho)\log\psi(\rho)}\frac{C_2(\rho)}{\sqrt{1-\rho^2}} \tag{1.90}$$

$$= \frac{-\rho + \sqrt{1-\rho^2}\,C_1(\rho) + \frac{\phi(\eta(\rho))}{-\psi(\rho)\log\psi(\rho)}\sqrt{1-\rho^2}\,C_2(\rho)}{1-\rho^2} \tag{1.91}$$

Suppose first that $\psi(\rho) \to c_0 > 0$ as $\rho \to \rho^*$, so that $-\psi(\rho)\log\psi(\rho)$ tends to a finite limit, since $\psi < 1$. Then all the terms in eq. (1.90) tend to a finite value except for $C_1(\rho)/\sqrt{1-\rho^2}$, which tends to ∞ as $\rho \to \rho^*$, so that $\tilde{\xi}'(\rho) \to \infty$ as well, so that $\xi(\rho)$ decreases towards 0 as ρ decreases towards ρ^*.

On the other hand, if $\psi(\rho) \to c_0 \leq 0$, then there are two cases to consider: since, $\sqrt{1-\lambda^2}\rho - \sqrt{1-\rho^2}$ in the numerator of C_2 is positive and increasing for all $\rho > \rho^*$, it can either be the case that $C_2 \geq 0$ for all $\rho > \rho^*$, or $C_2 < 0$ for all $\rho < \underline{\rho}$ for some $\underline{\rho} > \rho^*$. In the former case, the last term in eq. (1.90) is non-negative so that the by the same argument as above, $\tilde{\xi}'(\rho)$ is positive as $\rho \to \rho^*$. In the latter case, $\psi(\rho)$ is decreasing for all $\rho^* < \rho < \underline{\rho}$, but since $\psi(\rho) \to c_0 \leq 0$, it must be that $\xi(\rho) = 0$ for all ρ in this range. Hence, $\xi(\rho) \to 0$ in a weakly monotonic manner.

Analogously, suppose first that $\psi(\rho) \to c_0 > 0$ as $\rho \to 1$, so that $-\psi(\rho)\log\psi(\rho)$ tends to a finite limit. Since also C_1, $\phi(\eta(\rho))$ and C_2 tend to finite limits, all terms except $-\rho$ in the numerator of eq. (1.91) vanish, so that $\tilde{\xi}'(\rho) \to -\infty$ and $\xi(\rho)$ decreases towards 0 as ρ increases towards 1.

We now need check that this holds even if $\psi(\rho) \to c_0 \leq 0$ as $\rho \to 1$. Again, there are two possibilities depending on whether $\psi(\rho)$ is decreasing or increasing as towards c_0. In the former case, $\psi'(\rho) \to -\infty$, which only strengthens the fact that $\xi'(\rho) \to -\infty$ as $\rho \to 1$. In the latter case, by a similar logic to the one used above for $\rho \to \rho^*$, if $\psi(\rho) \to c_0 \leq 0$ and $\psi(\rho) > 0$ as $\rho \to 1$, it must be that $\xi(\rho) = 0$ as $\rho \to 1$, so that $\xi(\rho) \to 0$ in a weakly monotonic manner. ■

Chapter 2

Custodes invicem custodiunt: Commitment through competition by specialists in violence

2.1 Introduction

The enforcement of property rights and contractual agreements ultimately depends on the presence of agents, such as the police or the military, who can use coercive power to punish those who violate them. But how can these agents commit not to abuse this power for their own gain? This commitment is important in modern economies where the possibility of ex-post expropriation would seriously undermine incentives for ex-ante investments leading to bad economic outcomes[1], but where the means of coercive power are solely in the hands of specialized agents whom we call *specialists in violence* following the terminology of North, Wallis, and Weingast (2005). Thus, modern societies have agents whose job it is to guard property rights and contractual agreements, but "who guards the guards themselves? (*quis custodiet ipsos custodes?*)", as the famous question goes.

Our answer to this question is that "the guards guard each other" (*custodes invicem custodiunt*), that is, competition between specialists in violence, and in particular, their inability to commit not to turn against one another, keeps predatory behaviour at bay. In our model, even if specialists in violence could expropriate all output costlessly, it is attractive to protect producers from predators. This is because there is a marginal defensive advantage and consequently defence is an effective way to potentially eliminate other specialists in violence, reducing competition and leading to higher future payoffs. Producers can therefore engineer a Prisoner's dilemma that exploits the desire of specialists in violence to get rid of competitors, to threaten potential predators with elimination.

To illustrate the basic insight of our model more concretely, suppose there are two generals, commanding equally powerful armies, with no external threats. If they both decide to predate they take all output and keep half each. If they both decide to defend then they are paid a transfer, which we can think of as a tax or salary or even protection money, by the producers

[1] See Besley and Ghatak (2010) for an overview of links between expropriation and economic outcomes.

and do nothing. But if one of them defends and the other predates, then producers help the defending general fight against the predating one so that the probabilities of victory are greater than and less than half, respectively. If the defender wins then he will be the sole general left, so that he will be able to take all output for himself. Whoever loses the fight gets nothing. In this game, when the other general is predating, the payoff from defence consists of output times the probability of winning, which is greater than a half due to the producers' help. On the other hand, the payoff from colluding with the predating general is only half of output since they share output equally. Then producers can avoid predation by offering a transfer that makes each general prefer taking that transfer and doing nothing to being a predator fighting against the other general. This is how competition between the two generals lowers the level of expropriation.

By extending this logic to the case of many specialists in violence, we show that the proportion of output that they obtain in the form of transfers is decreasing in their number. Our model easily accommodates heterogeneity in strength among specialists in violence and we show how the level of expropriation is decreasing as the distribution of strength becomes more equal. Our paper makes the point that increasing competition between specialists in violence, both by increasing their numbers and making their strengths more equal is beneficial to producers, which is in line with the intuition that making power more diffuse reduces the incentives to abuse it. We also generalise the model in a different direction by allowing heterogeneity in the loyalty that specialists in violence command over their troops.

Finally using only within country variation over time we find that the positive effect of competition among specialists in violence on expropriation risk that we highlight in the model holds true for countries at lower levels of development but attenuates at higher levels of development. This is consistent with the idea that problem of civilian control over specialists in violence is a salient issue for countries at a less advanced stage of institutional development.

Our paper contributes to the large literature in economics and political science that attempts to explain the existence of the commitment by those who have power to expropriate from those who don't. The main thrust of the existing literature is that commitment arises as a consequence of the repeated nature of the game that producers and specialists in violence play. In a one-shot game producers anticipate predation at the end of the period and this leads to no investment in equilibrium.[2] But if this interaction is repeated infinitely, producers can play trigger strategies that make it attractive for specialists in violence to forgo predation in the present in exchange for larger payoff in the future. For this mechanism to sustain commitment, it is necessary that agents have a high enough discount factor, i.e., that they care enough about future payoffs. In this setup, competition between specialists in violence can be detrimental to to economic incentives as it can reduce their survival probability and hence the value of future output. Olson (1993) famously couched this view in terms of "roving bandits" whose precarious survival leads to full predation versus a "stationary bandit", an entrenched monarch who enjoys a long time horizon.[3]

[2]It is interesting to note that the problem of commitment becomes salient only in economies where output depends on ex ante investment. In a pure exchange economy the ability to commit is irrelevant since the equilibrium is likely to be Pareto efficient even with predation since there are no incentive effects. Piccione and Rubinstein (2007) present a model that makes this point formally.

[3]This argument is made formally in McGuire and Olson (1996) and Grossman and Noh (1990).

Our paper is inspired by the fact that some real world institutional arrangements seem at odds with this Olsonian view and are predicated on the often-voiced view that diffusion of power is good. For example, in order to avoid collusion leading to abuses of their power, there are often strict protocols governing the manner in which the highest ranks of the military meet. Another famous historical example, which we deal with in more detail later, comes from the Roman Republic, where ultimate power over the army was typically vested in two consuls with a view to keep a check on their power. This idea of checks and balances lies at the heart of our model, where the presence of several specialists in violence keeps each one in check creating a balance of power conducive to investments.

Our paper is related to Besley and Robinson (2009), who model the interaction between the military and civilian government when there is the possibility of the former seizing power through a coup. In their model, a key concern is the ability of the *government* to commit to pay the military, whereas our focus is on the commitment of the *military*. Furthermore, a major difference is that in our model specialists in violence can collude to expropriate fully without incurring any costs.

More broadly, our research agenda is similar to Acemoglu and Robinson (2006), but with the major difference that commitment arises not from the power of a specialist in violence to tie his own hands but from the existence of other specialists in violence who would stand to gain by punishing the deviant predator. This formulation enables us to attempt an answer to the question posed by Acemoglu (2003) about how specialists in violence can commit when the existence of their power to predate undermines any promises they make not to renege on their commitment whenever it is convenient. The insight that we formalise here is that commitment should not be seen as an additional strategy that may or may not be available to specialists in violence as a result of exogenous institutional arrangements. Instead, we argue that commitment should be seen as a feature of an equilibrium arising from a game played between more than one specialist in violence.

Our paper is also related to the large literature on the co-existence of economic activity and conflict.[4] This literature models choices of agents when agents can invest to produce as well as increase their predatory capacity. Typically some investment occurs even though this is lower than the first best where agents can commit not to predate. This literature assumes that all agents work as producers as well as specialists in violence or that within a unit where agents specialise, the producers and specialists in violence have solved their commitment problem. The key innovation that distinguishes our paper from this literature is that we attempt to unpack how commitment between producers and specialists in violence can arise in the first place.

The mechanism at play in our model is reminiscent of Dal Bó (2007), where a lobbyist can affect the outcome of a vote by a committee by offering members transfers which compensate voters for voting against their own preferences only when they are pivotal. Since this makes voting according to the wishes of the lobbyist a dominant strategy, the compensatory transfers are never paid out. The analogue idea in our model is that producers need to pay the specialists in violence only their payoff when they are the sole predator fighting against all others, i.e., when they are pivotal in predation, making this “bribe” small. On the other hand, our paper

[4]Examples include Skaperdas (1992), Hirshleifer (1995), and Grossman and Kim (1995). See Garfinkel and Skaperdas (2007b) for a survey of this literature.

does not assume the existence of any kind of contract enforcement, which is required in Dal Bó (2007).

Acemoglu, Egorov, and Sonin (2009) is another paper which incorporates some aspects of our model, in that it features elimination (through voting, rather than fighting) of competitors that can potentially improve future conditions, but their objective is to analyse what are stable configurations of power where no one can eliminate anyone else. In their context, in our model, any collection of specialists in violence is stable, since any predatory activity (including attacking others) will be punished by the other specialists who obtain the help of producers.

The paper is structured as follows. Section 2.2 discusses the baseline model with homogeneous agents and derives the comparative statics of the equilibrium. Section 2.3 extends the baseline model to allow heterogeneity in the strength of each specialist in violence. Section 2.4 extends the baseline model by introducing heterogeneity in the loyalty that specialists in violence command. Section 2.5 is a case study of the historical instution of consulship during the Roman Republic, which supports the intuition of our argument. Section 2.6 discusses our empirical results and Section 2.7 provides concluding remarks.

2.2 Model

The economy is populated by an exogenously given number of producers and specialists in violence. Producers operate a technology that requires some ex-ante investment in order to generate output. Specialists in violence, henceforth abbreviated to SIVs, can fight against each other and/or expropriate the producers' output. Specialisation is complete, so that producers cannot defend themselves against SIVs, whilst the latter cannot control the former's investment decisions. The interaction between these two groups is modelled as a game that unfolds as follow.

1. Producers make investments, whilst SIVs wait.
2. Output is realised and producers choose a fraction t of total output to offer to each of the SIVs.
3. Each SIV independently chooses whether to predate or defend.
4. (a) If all SIVs choose to defend then each is paid the transfer t by the producers and the game ends.
   (b) If some SIVs choose to predate, there is a fight between predators and defenders, with defeated SIVs obtaining a payoff of 0.
5. (a) If the predators win, they expropriate all output and share it equally among themselves, since producers cannot fight back.
   (b) If the defenders win, they enter a subgame where they are the only SIVs playing the same game, and producers once again make transfers and the game restarts from stage 3.

We first model the predation stage (the last three steps in the above timing) where SIVs make the decision of predating or defending. This decision depends on the transfers that are on offer from the producers. We then go back one step and derive the transfer that producers offer each SIV. After this, we model the stage where producers make ex-ante investments.

2.2.1 Fighting

Suppose that at this stage, $p > 0$ SIVs have decided to predate and $q > 0$ SIVs have decided to defend. The probability that the predators win is

$$\frac{p}{\delta q + p}, \tag{2.1}$$

whereas the probability that the defenders win is

$$\frac{\delta q}{\delta q + p}. \tag{2.2}$$

These probabilities are similar to those given by contest success functions commonly used in the conflict literature, but differ from the latter since they depend solely on the number of agents on each side of the fight and not on the effort exerted by them. Therefore, fighting is completely costless in this formulation.[5]

The parameter δ indicates the degree by which the technology of fighting favours defenders and we will make use of the following assumption regarding it.

Assumption 2.1 Defending SIVs have a combat advantage over predators, so that $\delta > 1$.

This assumption is foundational to our results. We can think of the defensive advantage as arising out of the possibility that producers help defending SIVs in the fight against the predating ones. Although in our model producers possess no combat ability, they could still provide help to defending SIVs through non-armed resistance in the form of intelligence gathering, sabotage or strikes, etc. Such activities would be of limited use to producers in protecting themselves from expropriation but could be a boost to a military force that can take advantage of them. Alternatively we can also think of the defensive advantage as arising from the possibility that troops of a SIV are more likely to obey a command to protect the producers rather than a command to predate. Although a defensive advantage is crucial in our model, it should be noted that this advantage can be arbitrarily small.[6]

2.2.2 Predation vs defence

Since by this stage output is already realised, we will normalize it to 1, so that all payoffs are fractions of total output. Consider a SIV's decision to predate or defend when there are p predators and q defenders. If he joins the predators, their number increases to $p + 1$ so that the probability of them winning is $\frac{p+1}{\delta q+p+1}$. Should they successfully predate, each SIV would

[5]Introducing an exogenous cost to conflict in this framework is straightforward and only strengthens our result further, since the outside option to co-operation with producers becomes less attractive. On the other hand, introducing endogenous fighting costs when there are multiple SIVs is not quite as straightforward, since the usual contest function approach cannot be easily extended to the case with many players divided into two factions.

[6]Note that an alternative way of specifying these probabilities for predators and defenders is $\frac{(1-\gamma)p}{(1-\gamma)p+\gamma q}$ and $\frac{\gamma q}{(1-\gamma)p+gq}$ respectively. This is equivalent to our formulation. The assumption analogous to assumption 2.1 that would ensure a defensive advantage would be $\gamma > 1/2$.

obtain a share $\frac{1}{p+1}$ of output, so that the expected payoff from joining p predators is

$$\Pi_q^{p+1} := \frac{1}{\delta q + p + 1} . \tag{2.3}$$

Should he instead join the defenders, their number rises to $q+1$ so that the probability of the defenders winning is $\frac{\delta(q+1)}{\delta(q+1)+p}$. After a successful defence, the remaining SIVs enter a subgame where they are offered transfers by producers and then choose to predate or defend. In that subgame, a SIV has the option of predating and getting at least the payoff from being the sole predator.[7] Then, the expected payoff from joining q defenders is at least

$$\begin{aligned} \Delta_{q+1}^p &:= \frac{\delta(q+1)}{\delta(q+1)+p}\Pi_q^1 \\ &= \frac{\delta(q+1)}{\delta(q+1)+p}\frac{1}{\delta q+1} . \end{aligned} \tag{2.4}$$

Given these payoffs from predation and defence, the following lemma shows that the latter dominates the former.

Lemma 2.1 Iff $\delta > 1$, $\Delta_{q+1}^p \geq \Pi_q^{p+1}$ for all p and q, with strict inequality if $p > 0$.

PROOF. Comparing Δ_{q+1}^p and Π_q^{p+1} we have

$$\begin{aligned} \frac{\delta(q+1)}{(\delta(q+1)+p)(\delta q+1)} &\geq \frac{1}{\delta q+p+1} \\ \iff \quad \frac{\delta q+p+1}{\delta q+1} &\geq 1 + \frac{p}{\delta(q+1)} \\ \iff \quad p\delta(q+1) &\geq p(\delta q+1) \end{aligned}$$

iff $\delta > 1$, with strict inequality if $p > 0$. □

This lemma shows that when there is a defensive advantage, a SIV strictly prefers to join forces with defending SIVs rather than the predators, if there are any of the latter. This is because the payoff from defending first and predating in the subsequent subgame, where some SIVs have been eliminated, is strictly greater than the payoff from predation. This means that in every subgame, there will be at most one predator.

2.2.3 Transfers

In the last stage, we saw that, from the point of view of an individual SIV, it is always better to defend than to predate if some of the other SIVs are predating. But what about when all the other SIVs are also defending? In that case, the transfers that the producers offer will determine the choice of whether to predate or defend.

In our model, producers make a take-it-or-leave-it offer to the SIVs, who then independently decide their actions. Then, given that producers have all the bargaining power, it follows that

[7] Note that for fixed $p+q$, Π_q^{p+1} is increasing in p.

SIVS are always pushed to their outside option.[8] This means that in every subgame after a successful defence, the producers' transfer is exactly equal to an individual SIVS payoff from becoming the sole predator, so that Δ^p_{q+1} as defined in (2.4) is the actual defence payoff, not merely its lower bound. Since this makes SIVS indifferent between being sole predators and defenders we will make the following assumption.

Assumption 2.2 SIVS who are indifferent between predating and defending choose defence.

We make defence the preferred option in case of indifference in order to rule out equilibria where only one SIV predates and everyone (including the producers) gets exactly the same expected payoff as in the case where all SIVS accept the producers' offer. [9] However such equilibria are purely an artifact of the producers pushing the SIVS to their outside option, and disappear as soon as the latter have some bargaining power. Given this assumption, the preceding arguments lead to the following proposition.

Proposition 2.1 The unique subgame-perfect Nash equilibrium of the game with $s + 1$ SIVS consists of producers offering each SIV a fraction

$$t = \frac{1}{1 + \delta s} \tag{2.5}$$

of total output, with all SIVS choosing not to predate.

PROOF. The proof is established by induction on the number of SIVS. Firstly, note that when there is only one SIV, his expected payoff from predation is one, since that is the probability with which he avoids mutiny and becomes an actual predator. Then, producers can ensure that he does not predate by $t = 1$: this would make the SIV indifferent between predation and non-predation, and by assumption 2.1 the SIV would not predate.

Next, suppose that we have already managed to prove that the proposition holds whenever the number of SIVS is less than or equal to some number s, and let us examine whether the proposition still holds if there are $s + 1$ SIVS.

To analyse the predation and defence payoffs of an individual SIV, suppose that $p \geq 1$ of the other SIVS have decided to predate and $q \leq s - 1$ have decided to defend. Then his payoff from joining the p other predators is

$$\frac{p+1}{p+1+\delta q}\frac{1}{(p+1)} = \Pi^{p+1}_q . \tag{2.6}$$

On the other hand, the payoff from joining the q defenders is the expected value of the product of the probability that $q + 1$ defenders win against p predators and of the payoff in the subgame where the defenders have won and there are only $q + 1$ remaining SIVS. Since

[8]The results are robust to changing the bargaining power of the producers and SIVS as long as SIVS do not have all the bargaining power. With full bargaining power SIVS make a take it or leave it offer leaving producers with nothing and consequently the incentive for ex-ante investment is destroyed.

[9]The only difference with these equilibria is that unlike the unique equilibrium in proposition 2.1 with no predation, these contain a positive probability of predation. However the expected level of expropriation is equal to the total transfers in the no predation equilibrium and moreover the central message of the paper about decrease in expropriation through increased competition remains a feature of these equilibria.

we are considering subgame-perfect equilibria we now that the payoff in that subgame will be the Nash equilibrium of that subgame. Furthermore, we assumed that the proposition holds in any game where the number of SIVs is at most s so that the Nash equilibrium payoff in a subgame where there are only $q+1$ SIVs is $\frac{1}{1+\delta q}$. The payoff from defence is then

$$\frac{\delta(q+1)}{p+\delta(q+1)}\frac{1}{1+\delta q} = \Delta^{p}_{q+1} \tag{2.7}$$

By lemma 2.1, $\Delta^{p}_{q+1} > \Pi^{p+1}_{q}$ for all values of p, with strict inequality since $p \geq 1$. Therefore a SIV always strictly prefers defence to predation if there is at least one other potential predator.

Suppose instead that, from the point of view of an individual SIV all of the other SIVs are defenders. Then his payoff from predation is $\frac{1}{\delta s+1}$, whereas that from defence is simply the transfer t. By assumption 2.2, producers can ensure that this SIV does not predate by offering a transfer exactly equal to his predation payoff. Therefore, when there are $s+1$ SIVs, the only equilibrium is one where producers offer $t = \frac{1}{\delta s+1}$ and all SIVs do not predate. □

To reiterate, the intuition of this result is as follows. Although a larger number of predating SIVs increases the probability of a successful predation, the payoff conditional on success is weighed down by the decreased share each SIV receives.[10] As a result it is more attractive for a SIV to stave off predation with the expectation of the larger share he receives if the defenders win. Even a marginal defensive advantage ensures that it is a dominant strategy for all SIVs to defend. If all other s SIVs are defending the payoff of a lone SIV who considers predation is $\Pi^{1}_{s} = \frac{1}{1+\delta s}$. Hence when producers offer him this amount they make him indifferent between predation and defence and given assumption 2.2, he defends.

It is convenient to define the expropriation rate that the producers face, i.e., the fraction of total output that they transfer to the SIVs as

$$\tau := (s+1)t = \frac{s+1}{1+\delta s}. \tag{2.8}$$

Taking the derivative of τ with respect to s we find that

$$\frac{\partial \tau}{\partial s} = -\frac{\delta - 1}{(1+\delta s)^2} < 0\,, \tag{2.9}$$

since $\delta > 1$ by assumption 2.1. This shows that not only is the transfer paid to an individual SIV decreasing in s, but that the sum of transfers is also decreasing in the number of SIVs. This is because, as the number of SIVs increases, the deviation payoff from predation becomes worse, which in turn decreases the equilibrium transfer paid to SIVs.

Remark 2.1 Expropriation is decreasing in the number of specialists in violence.

[10] It is interesting to note that the reason why the increase in the numerator of the probability of successful predation is exactly offset by the reduction in the share of each SIV is because p enters linearly in the numerator of the probability of successful predation defined in equation (2.1). Allowing for a more general functional form $\frac{f(p)}{\delta f(q)+f(p)}$ changes the results. Typically the uniqueness of equilibrium may no longer be available with a general $f(p)$ as multiple stable coalitions between SIVs may arise.

This result captures the mechanism that this paper highlights. Total expropriation tends to decrease when power is diffuse. In particular, total expropriation decreases in the number of SIVs as the balance of power between them is such that unilateral predation becomes more and more unattractive. This result is interesting when contrasted with the Olsonian idea that decreasing the number of SIVs decreases their incentives to expropriate fully. The two mechanisms may be seen as complementary to one another; it is possible to imagine that the number of SIVs arises at a point where these two forces equilibrate one another.

As we would expect, total expropriation is decreasing in the defensive advantage. The intuition for this is straightforward. As defence becomes easier, the expected payoff from predation decreases. Consequently SIVs are satisfied with a lower transfer and the tax rate the producers face goes down.

The central message of the model is that competition among specialists in violence creates a balance of power that makes predation unattractive, leading to a commitment not to predate. The intuition behind this result is simple: the defensive advantage not only skews the probability of combat victory towards defence, but makes it profitable to defend first and predate later, rather than predate at the outset; defence is a way to eliminate competitors and thus guarantee a bigger payoff for oneself, making it the dominant strategy. The inability to commit to refrain from using co-operation with producers as a way to get rid of each other places specialists in violence in a Prisoner's Dilemma, which the producers can exploit to avoid full predation.

The inability of specialists in violence to commit is a crucial issue in our paper. In societies like ours, the ability to commit to agreements arises precisely from the existence of agents who can use their specialisation in violence to punish those who renege on their commitments. But the commitment not to abuse their power is not available to the very agents who perform this enforcement function. Appealing to institutions to generate such commitment merely shifts the burden to the higher level specialists in violence who must support such institutions. This logic leads to an infinite regress where commitment at one level is sustained by commitment at a higher one. We have attempted to find a solution to this problem by using a somewhat different approach. In our model, what underlies the ability of specialists in violence to commit is not other institutions, but simply material forces that govern the success or failure of an attack aimed at expropriation, in other words material forces that shape the nature of the game that specialists in violence play.

2.3 Heterogeneity in strength

In this section we extend the model to allow SIVs to have differing strengths. This allows us to examine how expropriation changes in response to changes in the distribution of strength between SIVs . In particular we find that total expropriation decreases as the distribution of strengths becomes more equal. This strengthens our main point about the positive impact of competition between SIVs.

Suppose that the SIVs are indexed by i, where $i = 1, ..., s$, and let each SIV have strength x_i, which captures all factors that would contribute to increasing the probability of winning, such as the number of troops, their level of training or the quality of their equipment. fight between predators with total strength P and defenders with total strengths Q, the probability of victory are where we have rewritten equations (2.1) and (2.2) by dividing throughout by $1 - \delta$, so that $\delta = \delta/(1 - \delta)$. The assumption that $\delta > 1/2$ now corresponds to $\delta > 1$. Now that strengths are different, it is natural to assume that victorious predators share output proportionally to their strengths. Thus a SIV with strength x who successfully predated with other SIVs who have total strength P, would get a share of $\frac{x}{x+P}$ of total output.

We next prove the counterpart to lemma 2.1, showing that defence is a dominant strategy, being strictly dominant if there is at least one predator already.

Lemma 2.2 Iff $\delta > 1$, $x > 0$,

$$\frac{\delta(Q+x)}{P+\delta(Q+x)} \frac{x}{x+\delta Q} \geq \frac{P+x}{P+x+\delta Q} \frac{x}{x+P} \tag{2.10}$$

with strict inequality if $P > 0$.

Proof. Inequality (2.10) is true iff

$$\frac{\delta(Q+x)}{P+\delta(Q+x)} \frac{1}{x+\delta Q} - \frac{1}{P+x+\delta Q} \geq 0 \tag{2.11}$$

$$\frac{(\delta-1)Px}{(P+\delta(Q+x)(x+\delta Q)(P+x+\delta Q)} \geq 0\,, \tag{2.12}$$

which holds with strict inequality iff $\delta > 1$. □

We can now prove the analogue of proposition 2.1.

Proposition 2.2 The unique subgame-perfect Nash equilibrium of the game where each SIV has strength x_i is for producers to offer to each SIV a transfer

$$t_i^* = \frac{x_i}{x_i + \delta \sum_{j \neq i} x_j}\,, \tag{2.13}$$

and for all SIVs to not predate.

Proof. The proof is the same as that for proposition 2.1 but using lemma 2.2 to establish that defence is a dominant strategy whenever there is at least one predator, so that producers only need to offer to each SIV their payoff from being the sole predator. □

An interesting feature of the equilibrium is that each SIV's payoff depends not only on his strength, but also on that of all others. It is then natural to ask how the distribution of strengths affects the total amount of output that producers end up giving to the SIVs. The following proposition shows that a more equal distribution leads to lower transfers.

Proposition 2.3 Suppose that SIVs i and j have strengths $x_i > x_j$. Then reducing i's strength to $x_i - \varepsilon$ and increasing j's to $x_j + \varepsilon$, where $0 < \varepsilon < x_i - x_j$, will reduce total transfers paid to SIVs.

PROOF. Since the redistribution of strength keeps the sum of i and j's strengths constant, the payoff to all other SIVs is unaffected. Therefore, it suffices to show that the transfers to i and j, namely $t_i^* + t_j^*$, will fall. Then we need to show that

$$\begin{aligned}
&\frac{x_i}{x_i + \delta x_j + \delta \sum_{k \neq i,j} x_k} + \frac{x_j}{x_j + \delta x_i + \delta \sum_{k \neq i,j} x_k} \\
&\quad \geq \frac{x_i - \varepsilon}{x_i - \varepsilon + \delta(x_j + \varepsilon) + \delta \sum_{k \neq i,j} x_k} + \frac{x_j + \varepsilon}{x_j + \varepsilon + \delta(x_i - \varepsilon) + \delta \sum_{k \neq i,j} x_k} \\
&\quad = \frac{x_i - \varepsilon}{x_i + \delta x_j + (\delta - 1)\varepsilon + \delta \sum_{k \neq i,j} x_k} + \frac{x_j + \varepsilon}{x_j + \delta x_i - (\delta - 1)\varepsilon + \delta \sum_{k \neq i,j} x_k}. \qquad (2.14)
\end{aligned}$$

Letting $\sigma_i = x_i + \delta x_j + \delta \sum_{k \neq i,j} x_k$ and $\sigma_j = x_j + \delta x_i + \delta \sum_{k \neq i,j} x_k$, we need to show that

$$\frac{x_i}{\sigma_i} + \frac{x_j}{\sigma_j} = \frac{x_i \sigma_j + x_j \sigma_i}{\sigma_i \sigma_j} \qquad (2.15)$$

$$\geq \frac{x_i - \varepsilon}{\sigma_i + (\delta - 1)\varepsilon} + \frac{x_j + \varepsilon}{\sigma_j - (\delta - 1)\varepsilon} \qquad (2.16)$$

$$= \frac{x_i \sigma_j + x_j \sigma_i - 2(\delta - 1)\varepsilon\big(x_i - x_j - \varepsilon\big)}{\sigma_i \sigma_j + (\delta - 1)^2 \varepsilon\big(x_i - x_j - \varepsilon\big)}, \qquad (2.17)$$

which is true if $\delta > 1$ and $0 < \varepsilon < x_i - x_j$. □

This proposition shows that a Dalton-transfer of strength from a stronger SIV to a weaker one will reduce total transfers. As a consequence, a more equal distribution of strengths yields lower total transfers to SIVs, with the minimum being achieved when all SIVs are homogeneous.

Remark 2.2 Expropriation decreases with more equal distribution of strength among specialists in violence.

This is in line with the intuitive idea that a balance of power as arising from power being equally spread out over a number of agents helps in preventing predation. A more even distribution of power yields more effective competition, strengthening our main point that competition is the force underlying the ability of SIVs to commit. Seen together remarks 2.1 and 2.2 reinforce the positive impact that competition among specialists in violence has on investment incentives in the economy.

2.4 Heterogeneity in loyalty

In this section, we enrich the model in a different direction by endowing each SIV with a type parameter θ that represents the ability of a SIV to induce his troops to follow an order that would be considered predatory, i.e., an order to expropriate output from producers. We add one stage to the game where, immediately after the SIVs have decided whether to predate or defend, each potential predator independently experiences a revolt or mutiny by his troops with probability $1 - \theta$; in such an eventuality the SIV obtains a payoff of 0 and plays no further part in the game. We will call a SIV whose troops obey his order to predate an *actual* predator;

thus a potential predator becomes an actual predator with probability θ. Defenders do not face any possibility of mutiny.

We can interpret the average level of θ across SIVs embodies institutional factors such as the professionalism of troops or the extent to which troops respect civilian law: a society where the military is composed of professional soldiers is likely to have a higher average θ than one where defence is provided by popular militias, since soldiers who are also producers themselves would be more likely to disobey a command to expropriate output; similarly, generals who command troops drawn from a society where there is a strict separation between military and civilian spheres would have on average lower θs than generals whose troops are regularly used to intervening in civilian matters.

On the other hand, variation across SIVs in θ captures individual differences in leadership and charisma between leaders: an example of a high-θ SIV would be Julius Caesar who, knowing the loyalty of his veterans from the Gallic Wars, crossed the Rubicon and marched upon Rome in open defiance of a senatorial command.

The natural question that arises when SIVs are heterogeneous is whether their type θ is observable or not. In the following subsection we analyse the case where θ is observable by all agents, including producers. Subsequently we explore what happens in case where a SIV's θ is his own private information. For simplicity, we no longer allow agents to make occupational choices making $s + 1$, the number of SIVs, exogenously determined. Also, we will normalise output to 1 and think of all payoffs as fractions of total output.

2.4.1 Complete Information

It is easy to see that the baseline model corresponds to the case where all SIVs have their θs equal to 1, so that there is no possibility of mutiny. It turns out that when we generalise the baseline model to admit heterogeneity in θ that proposition 2.1 extends in a very natural way to give us the following proposition.

Proposition 2.4 The unique subgame-perfect Nash equilibrium of the complete information game with $s + 1$ SIVs consists of the producers offering each SIV a fraction

$$t^*_{s+1}(\theta) = \theta \frac{1 - \delta}{(1 - \delta) + \delta s} \tag{2.18}$$

of total output, where θ is an individual SIV's type, with all SIVs choosing not to predate.

PROOF. The underlying logic of the proof is the same as the arguments leading up to proposition 2.1, but we prove it formally by induction on the number of SIVs.

Firstly, note that when there is only one SIV with type θ, his expected payoff from predation is θ, since that is the probability with which he avoids mutiny and becomes an actual predator. Then, producers can ensure that he does not predate by $t^* = \theta$: this would make the SIV indifferent between predation and non-predation, and by assumption 2.1 the SIV would not predate.

Next, suppose that we have already managed to prove that the proposition holds whenever the number of SIVs is less than or equal to some number s, and let us examine whether the proposition still holds if there are $s+1$ SIVs.

To analyse the predation and defence payoffs of an individual SIV with type θ, suppose that $p \geq 1$ of the other SIVs have decided to predate and $q \leq s-1$ have decided to defend. Since this SIV must decide whether to predate or defend before he observes whether the other p potential predators' troops mutiny or not, the number of other actual predators is a random variable P, ranging from 0 to p. Then his payoff from joining the p other potential predators is

$$\theta \mathbb{E}\left(\frac{(1-\delta)(P+1)}{(1-\delta)(P+1)+\delta q} \frac{1}{P+1}\right) = \theta \mathbb{E}\left(\Pi_q^{P+1}\right) . \tag{2.19}$$

On the other hand, the payoff from joining the q defenders is the expected value of the product of the probability that $q+1$ defenders win against P actual predators and of the payoff in the subgame where the defenders have won and there are only $q+1$ remaining SIVs. Since we are considering subgame-perfect equilibria we now that the payoff in that subgame will be the Nash equilibrium of that subgame. Furthermore, we assumed that the proposition holds in any game where the number of SIVs is at most s so that the Nash equilibrium payoff in a subgame where there are only $q+1$ SIVs is $\theta \frac{1-\delta}{(1-\delta)+\delta q}$. The payoff from defence is then

$$\mathbb{E}\left(\frac{\delta(q+1)}{(1-\delta)P+\delta(q+1)} \theta \frac{1-\delta}{(1-\delta)+\delta q}\right) = \theta \mathbb{E}\left(\Delta_{q+1}^P\right) \tag{2.20}$$

By lemma 2.1, $\Delta_{q+1}^P \geq \Pi_q^{P+1}$ for all values of P, with strict inequality whenever $P > 0$, which happens with some probability since P ranges from 0 to $p \geq 1$. Therefore $\mathbb{E}\left(\Delta_{q+1}^P\right) > \mathbb{E}\left(\Pi_q^{P+1}\right)$, so that a SIV always strictly prefers defence to predation if there is at least one other potential predator.

Suppose instead that, from the point of view of an individual SIV with type θ, all of the other SIVs are defenders. Then his payoff from predation is $\theta \frac{1-\delta}{\delta s+(1-\delta)}$, whereas that from defence is simply the transfer t_{s+1}^*. By assumption 2.2, producers can ensure that this SIV does not predate by offering a transfer exactly equal to his predation payoff. Therefore, when there are $s+1$ SIVs, the only equilibrium is one where producers offer $t_{s+1}^*(\theta) = \frac{1-\delta}{\delta s+(1-\delta)}$ and all SIVs do not predate. □

It might be somewhat surprising at first to see that each SIV's transfer depends only on his own type and is independent of any other SIV's type. But this is a straightforward consequence of the logic that sustains the equilibrium: because the defensive advantage excludes the possibility of there being more than one predator, each SIV's alternative to accepting the trasfer is simply his payoff from being the sole predator when everyone else defends; since the type is irrelevant for defence, the defenders' type is irrelevant to this payoff.

On the other hand, the overall distribution of types becomes important if the type is private information of each SIV, as we discuss in the following section.

2.4.2 Incomplete Information

In this section we allow θ to take value θ_H with probability γ and θ_L with probability $1-\gamma$. Furthermore, the value of θ is private to each SIV, but is fully revealed after a fight between predators and defenders. Now the number of predators and defenders, denoted by P and Q, respectively, are random variables whose value depends on the realisations of the θs of the SIVs. This uncertainty means that, unlike the complete information case, conflict and expropriation can arise even in equilibrium, as the following proposition shows.

Proposition 2.5 Except when both $s = 0$ and $\theta_H = 1$, there exists a threshold γ^* such that if $\gamma \leq \gamma^*$, the unique symmetric subgame-perfect Bayesian equilibrium consists of the producers offering a transfer

$$\underline{t}^*_{s+1} = \theta_L \Pi^1_s - \left(\theta_L \mathbb{E}\left(\Delta^P_{Q+1} - \Pi^{P+1}_Q\right)\right) / \Pr(Q = s) \tag{2.21}$$

that makes low-θ SIVs indifferent between predating and defending, with high-θ SIVs preferring predation. In this case, the total payout by producers is

$$\underline{\tau}_{s+1} = (s+1)\mathbb{E}(\theta)\,\mathbb{E}\left(\Pi^{P+1}_Q\right) . \tag{2.22}$$

If $\gamma > \gamma^*$, the unique symmetric subgame-perfect Bayesian equilibrium consists of the producers offering a transfer $\overline{t}^*_{s+1} = \theta_H \Pi^1_s$ that makes high-θ SIVs indifferent between predating and defending, with low-θ SIVs strictly preferring defence. The total payout in this case is $\overline{\tau}_{s+1} = (s+1)\theta_H \Pi^1_s$.

PROOF. Given assumption 2.2, we only need to consider following four candidates for pure-strategy equilibria: (a) a pooling equilibrium where both types defend; (b) a pooling equilibrium where both types predate; (c) a separating equilibrium where high-type SIVs predate and low-type ones defend; (d) a separating one where low-type SIVs predate and high-type ones defend.

Consider first case (a), where both types defend. From the perspective of an individual SIV with type θ, taking as given the strategy of all other SIVs, his payoff from predation is identical to the complete information case, viz., Π^1_s, whereas his payoff from defence is simply the transfer that the producers make to the SIVs. By assumption 2.2, producers can avert predation by either type by offering the high-type's predation payoff, so that the equilibrium transfer would be $\overline{t}^*_{s+1} = \theta_H \Pi^1_s$. In this case, the total payment that producers make to SIVs is

$$\overline{\tau}_{s+1} = (s+1)\theta_H \Pi^1_s = \theta_H \frac{(1-\delta) + (1-\delta)s}{(1-\delta) + \delta s} < 1 \tag{2.23}$$

since $\delta > 1/2$ and we are excluding the case when both $s = 0$ and $\theta_H = 1$.

Consider next case (b) where both types predate. This cannot be an equilibrium because in this case the producers total payout would be 1; producers can always do better than this by offering the transfer in the pooling equilibrium of case (a).

Consider next the separating equilibrium of case (c) where high-types predate but low-types defend. The payoff from predation is then

$$\pi_P(\theta) = \mathbb{E}\left(\theta \Pi^{P+1}_Q\right) \tag{2.24}$$

since each SIV can be a predating high-type with probability γ and a defending low-type with probability $1-\gamma$. The payoff from defence is

$$\pi_D(\theta) = \Pr(Q=s)\big(\underline{t}_{s+1} - \theta\Pi_s^1\big) + \mathbb{E}\left(\theta\Delta_{Q+1}^P\right) \tag{2.25}$$

where $\underline{t}_{s+1}$ is the transfer made by producers should all s other SIVs be low-types. The payoff (2.25) follows from the fact that, since we are considering subgame-perfect equilibria, if a defender of type θ wins, his payoff will be the equilibrium payoff in a complete information game where all q other SIVs have type θ_L. By proposition 2.4 we know this is $t_{q+1}^*(\theta) = \theta\Pi_q^1$.

To find the lowest transfer $\underline{t}_{s+1}^*$ that ensures that low-types do not predate, we let $\pi_P(\theta_L) = \pi_D(\theta_L)$ and solve for $\underline{t}_{s+1}$ to obtain

$$\underline{t}_{s+1}^* = \theta_L\left(\Pi_s^1 - \frac{\mathbb{E}\left(\Delta_{Q+1}^P - \Pi_Q^{P+1}\right)}{\Pr(Q=s)}\right) \tag{2.26}$$

with $\mathbb{E}\left(\Delta_{Q+1}^P - \Pi_Q^{P+1}\right) > 0$ since $\Delta_{q+1}^p > \Pi_q^{p+1}$ for all q and for all $p>0$ (since we are excluding the case $s=0$) by lemma 2.1.

For $\underline{t}_{s+1}^*$ to be an equilibrium transfer, it must be high-types prefer predation to defence, i.e., $\pi_P(\theta_H) > \pi_D(\theta_H)$, which is equivalent to

$$\begin{aligned}\pi_P(\theta_H) - \pi_D(\theta_H) &= \Pr(Q=s)\left(\theta_H\Pi_s^1 - \underline{t}_{s+1}^*\right) + \theta_H\mathbb{E}\left(\Pi_Q^{P+1} - \Delta_{Q+1}^P\right) \\ &= \Pr(Q=s)(\theta_H-\theta_L)\left(\Pi_s^1 - \frac{\mathbb{E}\left(\Delta_{Q+1}^P - \Pi_Q^{P+1}\right)}{\Pr(Q=s)}\right) > 0\,.\end{aligned} \tag{2.27}$$

This means that a separating equilibrium exists if and only if γ and θ_H are such that $\underline{t}_{s+1}^*$ is strictly positive. Due to the technical nature of its proof, we show that the total payout in this separating equilibrium is $\underline{\tau}_{s+1} = (s+1)\mathbb{E}\left(\theta\Pi_Q^{P+1}\right)$ in lemma 2.A.4 in the appendix.

Finally, consider case (d) where low-types predate and high-types don't. If high-types prefer not predating it must be that $\pi_D(\theta_H) \geq \pi_P(\theta_L)$, which means that, by arguments analogous to those for (2.25),

$$\underline{t}_{s+1} \geq \theta_H\big(\Pi_s^1 - \mathbb{E}\left(\Delta_{Q+1}^P - \Pi_Q^{P+1}\right)\big) > \theta_L\big(\Pi_s^1 - \mathbb{E}\left(\Delta_{Q+1}^P - \Pi_Q^{P+1}\right)\big)\,. \tag{2.28}$$

This is equivalent to $\pi_D(\theta_L) > \pi_P(\theta_L)$ so that low-types too prefer not predating, resulting in a contradiction. Hence this cannot be an equilibrium.

To conclude the proof we need to show how the type of equilibrium, viz., pooling or separating, depends on γ. Suppose γ and θ_H are such that both the separating equilibrium in case (c) and the pooling equilibrium in case (a) exist. Then producers can choose the equilibrium that results in the lower payout by setting appropriate transfers. Since the payout $\overline{\tau}$ in the pooling equilibrium is independent of γ, whereas the payout $\underline{\tau}$ in the separating equilibrium is increasing in γ by lemma 2.A.4, there exists a threshold γ^* such that producers prefer the pooling equilibrium if $\gamma > \gamma^*$. □

2.4.3 Institutions and transitions

In addition to the prediction that expropriation decreases with the number of specialists in violence, this model differs from the Olsonian one in another way. In this model the payoff of specialists in violence is determined by their outside option of predation rather than their choice of the revenue maximising tax rate. This approach is similar in spirit to Acemoglu and Robinson (2006), where the payoff of rulers is determined by the constraints on their rule arising out of the possibility of a coup; this constraint is determined by the possibility of other generals replacing the king through a coup in the event he decides to increase his payoff through predation.

Monarchy can be mapped onto our model in a straightforward way by modelling the monarch as a specialist in violence who has a markedly higher θ than the others, i.e., his generals. This fits well with the idea that θ captures the likelihood of predatory orders being obeyed: historically, kings could draw their authority from special sources, such as divine right, imbuing their orders with a greater legitimacy than anyone else's. In such a regime we would expect the king to secure a high payoff relative to his generals and, as we show in proposition 2.4, this is indeed the case as the payoff of a specialist in violence is increasing in his θ.

Similarly, we can map the opposite Olsonian scenario of roving banditry. We could think of each bandit as being a separate specialist in violence, so that he retains full control over his own actions. This would be a natural interpretation of our baseline model where $\theta = 1$ for all specialists in violence.

In the region where $\gamma \leq \gamma^*$ we can see that the transfer paid to a defending SIV is decreasing in θ. We think of θ as a parameter that captures the ideological environment. If ideas of justice and fairness are firmly entrenched in an economy then θ, the probability with which troops obey orders for predation, would be low since troops would see such an order as being unjust and unfair. In medieval times the powers of the sovereign were believed to be of divine origin. Consequently the command of the sovereign was seen as being just in and of itself . The late seventeenth and early eighteenth century saw a change in the ideological environment. This can be seen in the arrival of social contract theorists such as John Locke and Jean-Jacques Rousseau who saw the idea of justice and fairness as being independent of the sovereign's command. Since this change in the ideological environment can be seen as a decrease in θ, it is possible to understand the reduction in payout to specialists in violence from renaissance onwards as being ideologically driven.

We have shown that the probability of conflict is non-monotonic in γ. At high levels of γ, in particular for $\gamma > \gamma^*$, the producers prefer the pooling equilibrium where they pay out a large transfer to all SIVs. In this region there is no conflict although output is quite low. One can think of this as a relatively stable although highly exploitative feudal regime. The incentives for investment in such a regime are quite low since producers correctly anticipate that a large part of their output would need to be paid out to SIVs.

When γ falls below γ^* the amount of conflict increases. This is the point at which the equilibrium shifts from being pooling to separating. The producers prefer paying out a lower transfer to defenders even though they face a positive probability of predation. In this region there would be frequent transitions where one set of SIVs fight the others and new regimes are established and overthrown. Note that although conflict increases, the expected payout by the producers is lower since the producer's payout is weakly increasing in γ. In a historical context, it is possible to think of this as the early modern period which was charaterised by frequent regime changes but at the same time lower expropriation than medieval times. Finally as γ decreases further the incidence of conflict decreases and a stage is reached where the expected payout by the producers is also quite low. It is possible to think of this as a modern stable state with high incentives for investment where the transfers paid to SIVs do not constitute a large fraction of the GDP.

What are the factors that cause γ to change? It is possible to think of γ as arising out of the occupational choices made by people who have the potential to be charismatic leaders. When technology of production is primitive, these agents have a low outside option in the productive sector. This causes them to become specialists in violence. As technology improves however, these agents move into production thereby reducing γ. The story that the model then delivers is consistent with European history over the last five hundred years where large technological progress and political evolution happened simultaneously. In future work it would be interesting to embed the incomplete information model in an occupational choice framework to see whether the model delivers this link from technological parameter α to the what would be an endogenous γ.

One of the long standing methodological debates in the discipline of history has been the one between the structuralist/Marxist view and the agency view. The former sees history as shaped by forces essentially out of the control of the individual agent. Marx for example sees the means of production as the driving force of history. As these change over time, changes in the political super-structure comes about. Interestingly, as Guiso, Sapienza, and Zingales (2006) points out, this view is also shared by the Chicago School. On the other hand the agency view states that individual agents such as Napolean, Robespierre, and Julius Ceasar do substantially alter history. In this view the actions of certain individual can lead history to a significant long run departure from its path in their absence. Our framework allows us to unify these two methodological positions. Material forces do affect the probability of a charismatic leader arising which would consequently lead to transitions and conflict, but at the same time the existence of these leaders is a necessary ingredient that shapes social dynamics.

2.5 Consuls in the Roman Republic

In this section we examine a particular institutional arrangement from ancient Rome that resonates quite cleanly with the mechanics of the model presented above. Consuls were the military and civil heads of the state during the Roman republic. The *fasti consulares*, a listing of the names and tenure of consuls, dates its first entry to 509 BC. The time period that fits our model most closely is from 509 BC when the office was established to around 89 BC.[11]

[11] A consul's power was superseded only in case of military emergency when a dictator was appointed. The instances of appointment of a dictator were few and short lived in this period. The exception to the rule of two consuls was the period of 426-367 BC which is known as 'the conflict of the orders' when consular power was often shared between three or more military tribunes. This does not affect our story since the results of our model are

Although the office of the two consuls persisted well after the establishment of imperial rule in Rome, the concentration of the *imperium* in two consuls, that is their status as the joint heads of the executive, diminished gradually once Sulla assumed dictatorial control in 89 BC. This decline continued under the appointment of Julius Caesar as a perpetual dictator in 44 BC and thereafter under the establishment of imperial rule under Augustus in 27 BC.

Two consuls were elected every year and jointly held the *imperium*. Any decision made by a consul, such as a declaration of war, was subject to veto by the other consul. As the military heads, consuls were expected to lead Roman armies in the event of a war. In case both consuls were in the battlefield at the same time, they would share the command of the army, alternating as the head on a day to day basis. The election of the consuls was held by an assembly of soldiers known as the *centuria*.[12] The fact that consuls were elected from within the military and by the military confirms the primacy of their role as the heads of military. Indeed, their roles as the civilian heads can be seen as arising from the control they wielded over the military. It is therefore appropriate to think of them as analogous to the specialists in violence in the model.

The crucial assumption that we make in the model is $\delta > 1$. This ensures that when the specialists in violence are evenly divided on both sides in a battle, the side supporting the producers has at least a marginal advantage. This assumption seems valid in this setting. During this period in Roman history, a potential soldier needed to prove ownership of a certain amount of property to be eligible for recruitment in the military. This meant that the soldiers tended to have close family who were typically engaged in productive activities such as agriculture. Consequently, if the two consuls disagreed on an order to predate, the military was at least marginally more likely to obey the order for protection of the producers over an order for predation. Knowing this both consuls would have preferred protecting the producers leading to the Prisoner's Dilemma that we highlight. It is interesting to note that the property requirement for recruitment into the army was finally relaxed in 107 BC. This was followed closely by the transition of the republic into a dictatorship first under Sulla in 89 BC followed later by Julius Caesar and eventually the establishment of a monarchy under Augustus in 27 BC.

This institutional arrangement points to the belief that two military heads would effectively balance each other out. Since together they enjoyed absolute power, there was nothing preventing them from colluding with each other, other than the architecture of the game itself. The possibility of collusion can arise either through infinite repetition of the one shot game or through the possibility of contracting. It is possible to identify the institutional features that precluded these. Yearly elections ensured a finite time horizon for the consuls. Consuls were barred from seeking re-elections immediately after serving a year in office. Usually a period of ten years was expected before they could seek the office again. This term limit preserved the

preserved as long as the number of specialists in violence is strictly greater than 1. We have relied on Hornblower and Spawforth (2003) as a reference for the historical material used in this case study.

[12] The assembly had 193 voting units, each unit representing a century, that is a group of one hundred soldiers. The assembly was composed of 18 centuries of *equites* that is the cavalry, 170 centuries of *pedites* that is the infantry and 5 centuries of non-combatants such as the horn blowers, artisans, etc. The voting order was the *equites* first followed by the *pedites* and lastly the non-combatants. See Taylor (2003) for a detailed exposition of the voting procedure in the *centuria*.

one-shot nature of the game. Second, there was no possibility of contracting since there was no higher authority than the consuls that could enforce any such contract. It appears that the consuls were locked in a game where the unique equilibrium was that they did not predate.

2.6 Empirical analysis

In this section we attempt to test part of our model. In particular we can test remark 2.1 that indicates that we should expect a negative relationship between the risk of expropriation and the number of SIVs. Unfortunately we don't have the data to test remark 2.2 which shows that the risk of expropriation is lower when the power of SIVs is more equal.

The empirical analysis is based on panel data on World Military Expenditures and Arms Transfers dataset compiled by the US Department of state.[13] The data comprises of 168 countries over an 11 year period from 1995-2005. This contains data on our main explanatory variable, the number of active troops per one thousand. It also contains data on military and government expenditure in 2005 US Dollars which we use as controls.

For our outcome variable we rely on the International Country Risk Guide (ICRG) compiled by Political Risk Services.[14] This contains an index that measures the risk of expropriation on a scale of 0-12 with a higher score indicating a lower risk. Our baseline specification is

$$y_{it} = \alpha_i + \beta_t + \gamma_1 \text{Armed Forces}_{it} + X_{it}'\lambda + \varepsilon_{it}. \tag{2.29}$$

The variable "Armed Forces" is the log of the number of active troops for one thousand people in the population. Note that the ideal empirical counterpart to SIVs is a variable that captures the number of military leaders who each command independent units. Since such data is unavailable we use the log of the number of armed forces instead. If the fraction of military leaders to armed forces remains constant within a country over the sample period, then there is no problem with using the armed forces variable. This is because the number of SIVs is some fraction θ_i of the number of armed forces. To see this mathematically note that

$$(1 - \tau_{it}) = c_i(\theta_i \cdot \#\text{ armed forces}_{it})^{\gamma_1} \tag{2.30}$$

$$\implies \ln(1 - \tau_{it}) = \ln c_i + \gamma_1 \ln \theta_i + \gamma_1 \ln(\#\text{ armed forces}_{it}) \tag{2.31}$$

The first two terms on the right hand side constitute the country fixed effect and cannot be identified separately. However the coefficient on the log of number of armed forces gives us an estimate of γ_1. The assumption underlying this is that the structure of military within a country, that is the proportion of soldiers and commanders stays constant.

As seen in remark 2.1, we should expect γ_1 to be positive. X_{it} is a vector of time varying country level controls including income as measured by log per capita GDP, log per capita government spending, log per capita military spending, log population. Since the risk of

[13] The data is available at http://www.state.gov/t/avc/rls/rpt/wmeat/2005/index.htm

[14] The investment profile component in the ICRG dataset has been widely used in the literature as a measure of risk of expropriation starting from Knack and Keefer (1995). As noted by Acemoglu, Johnson, and Robinson (2001), although the variable is designed to capture the risk of expropriation is for foreign investment, the correlation with the risk of expropriation for domestic investment is likely to be high.

expropriation and the proportion of population in the armed forces could also be correlated to levels of internal and external conflict, we control for these using indices for these two variables that are also part of the ICRG dataset. α_i and β_t are the country and time fixed effects.

Table 2.1 in the appendix reports the results of this regression. We observe that the estimate of γ_1 is close to zero and statistically insignificant in all specifications. The maintained hypothesis for this regression model is that the competition effect that we model applies equally to all countries. However it may be possible that the net effect of competition among the SIVs has a differential impact at different levels of development. In particular it is reasonable to believe that the threat of expropriation is real at lower levels of development when institutions are not well developed. On the other hand at advanced stages of institutional development, civilian control over the military is well established and consequently greater numbers within the armed forces ought not to affect the risk of expropriation. To test this hypothesis we regress the following specification where we allow the armed forces variable to interact with income

$$y_{it} = \alpha_i + \beta_t + \gamma_1 \text{Armed Forces}_{it} + \gamma_2 \text{Armed Forces}_{it} * \text{Income}_{it} + X'_{it}\lambda + \varepsilon_{it}. \quad (2.32)$$

Table 2.2 reports the results of this regression. We can see that now the estimate of γ_1 is positive and significant indicating that increasing the proportion of population in the armed forces reduces the risk of expropriation. Moreover the estimate of γ_2 indicates that as expected the competition effect is strong at low levels of development and attenuates with income.

We can also test this hypothesis by allowing the armed forces variable to have a differential impact if a country is a member of the OECD. We expect the coefficient on the interaction between OECD and armed forces to be negative since we don't expect competition among SIVs to affect the risk of expropriation within OECD countries. We run

$$y_{it} = \alpha_i + \beta_t + \gamma_1 \text{Armed Forces}_{it} + \gamma_2 \text{Armed Forces}_{it} * \text{OECD}_i + X'_{it}\lambda + \varepsilon_{it}. \quad (2.33)$$

Table 2.3 reports the results of this specification. Once again we observe that the estimate of γ_1 is positive and significant whereas the estimate of γ_2 is negative and significant. This indicates that the positive effect of competition among SIVs on investment incentives appears to be true for non OECD countries.

A potential concern with the 1995-2005 time period is that our results may be affected by the heterogeneous impact of the September 11, 2001 attacks. To address this we run our main specification from equation (2.32) on a sample restricted to 1995-2001. Table 2.4 reports the results. We observe that the results are not affected.

Another concern with these results is the endogeneity of variables such as current income, government and military expenditure, and conflict. We attempt to deal with this concern in two ways. First by taking an instrumental variables approach, and second by replacing contemporaneous regressors with their lags.

Our first attempt to address the endogeneity is through estimating the specification from equation (2.32) by using the lags of all variables on the right hand side. Table 2.5 reports the results. We see that the pattern of results continues to be the same as seen in table 2.2. γ_1 continues to be positive and significant whereas γ_2 continues to be negative and significant.

Table 2.6 reports the results from using the same set of instruments on the specification in equation (2.33). Once again we see the same pattern of results in relation to γ_1 and γ_2.

The instrumental variable approach is based on the identifying assumption that the lagged values of income, government expenditure, etc. do not have a direct impact on expropriation risk. Since this is unlikely to be entirely correct we also try using the lagged variables as regressors rather than as instruments. We run

$$y_{it} = \alpha_i + \beta_t + \gamma_1 \text{Armed Forces}_{it} + \gamma_2 \text{Armed Forces}_{it} * \text{Income}_{it-1} + X'_{it-1}\lambda + \varepsilon_{it} \quad (2.34)$$

where all the regressors except armed forces are lagged one period.[15] Table 2.7 reports the results of this regression. We see that although the magnitude of the effect drops, the result is consistent with the earlier specifications in that we find a positive and significant γ_1 and a negative and significant γ_2. Table 2.8 reports the results from regressing the lagged specification with the OECD indicator.

2.7 Conclusion

The ability to commit is one of the foundations of economic activity. This arises as a result of agents who specialise in enforcement of commitment through the threat of violence. How do these agents commit not to use their powers to expropriate others? This paper has attempted to answer this question. We have argued that commitment arises as an artifact of the Prisoner's Dilemma type game form within which these agents find themselves. Even though they could secure a higher payoff by colluding, they are unable to do so since unilateral adherence to their role as the protectors of the producers is always individually rational. Moreover our model shows how it is in the interest of the elite to have more diffuse power structure since that acts as credible commitment against abuse of power and as such is a first step towards a political Coase theorem.

Using within country variation to test the model, we find that competition among SIVs reduces the risk of expropriation, but only in developing countries. This is consistent with the fact that the problem of civilian control over SIVs is more salient at lower levels of institutional development. Our model therefore has implications about how to optimally structure the armed forces in less developed countries where civilian control over the military may be a problem.

[15] Since the model predicts a relationship between contemporaneous numbers in the armed forces and the risk of expropriation, we have not lagged the armed forces variable. However the results of the regression where the armed forces variable is also lagged one period are similar to the ones reported in table 2.7.

Appendix 2.A Mathematical appendix

Lemma 2.A.1 The distribution of P and Q when there are s SIVs, and where θ_H-SIVs predate and θ_L-SIVs defend is given by

$$\begin{aligned} \Pr(P = p, Q = q) &= \frac{s!}{p!q!(s-p-q)!}(\gamma\theta_H)^p(1-\gamma)^q(\gamma(1-\theta_H))^{s-p-q} && (2.35)\\ &= \binom{s}{q}\gamma^{s-q}(1-\gamma)^q\binom{s-q}{p}\theta_H^p(1-\theta_H)^{s-p-q}\,. && (2.36)\end{aligned}$$

so that the marginal distribution of Q is $\mathrm{Bin}(s, 1-\gamma)$ and the conditional distribution of P given Q is $\mathrm{Bin}(s-Q, \theta_H)$.

Lemma 2.A.2 If $X \sim \mathrm{Bin}(p, n+1)$,

$$\mathbb{E}\left(Xf(X)\right) = \mathbb{E}\left(X\right)\ \mathbb{E}\left(f(Y+1)\right)\ , \tag{2.37}$$

where $Y \sim \mathrm{Bin}(p, n)$.

PROOF. The expectation of $Xf(X)$ is

$$\begin{aligned} \mathbb{E}\left(Xf(X)\right) &= \sum_{x=0}^{n+1}\frac{(n+1)!}{x!(n+1-x)!}p^x(1-p)^{n+1-x}xf(x)\\ &= \sum_{x=1}^{n+1}\frac{(n+1)n!}{(x-1)!(n-(x-1))!}pp^{x-1}(1-p)^{n-(x-1)}f(x)\\ &= (n+1)p\sum_{y=0}^{n}\frac{n!}{y!(n-y)!}p^y(1-p)^{n-y}f(y+1)\\ &= \mathbb{E}\left(X\right)\ \mathbb{E}\left(f(Y+1)\right)\,. \end{aligned} \tag{2.38}$$

Lemma 2.A.3 If $X \sim \mathrm{Bin}(1-p, n)$ and f is a monotonically decreasing function, $\mathbb{E}\left(f(X)\right)$ is increasing in p.

PROOF. The derivative of $\mathbb{E}\left(f(X)\right)$ with respect to p is

$$\begin{aligned} \frac{d}{dp}\mathbb{E}\left(f(X)\right) &= \frac{d}{dp}\sum_{x=0}^{n}\frac{n!}{(n-x)!x!}p^{n-x}(1-p)^x f(x)\\ &= \sum_{x=0}^{n-1}\frac{n!}{(n-x-1)!x!}p^{n-x-1}(1-p)^x f(x)\\ &\quad -\sum_{x=1}^{n}\frac{n!}{(n-x)!(x-1)!}p^{n-x}(1-p)^{x-1}f(x)\\ &= n\sum_{x=0}^{n-1}\binom{n-1}{x}p^{n-1-x}(1-p)^x(f(x)-f(x+1)) > 0 \end{aligned} \tag{2.39}$$

if f is monotonically decreasing. □

Lemma 2.A.4 In a separating equilibrium where high-types predate and low-types defend, the producers' total payout is given by

$$\underline{t}_{s+1} = (s+1)(\gamma\theta_H + (1-\gamma)\theta_L)\mathbb{E}\left(\Pi_Q^{P+1}\right) \tag{2.40}$$

and this is increasing in γ.

PROOF. Let $\widehat{\Delta}_Q^P$ denote a random variable takes value Δ_Q^P is $Q < s+1$ and $\underline{t}_{s+1}$ if $Q = s+1$. Also, let P' and Q' denote the number of predators and defenders out of $s+1$ SIVs, rather than s SIVs as for P and Q, and let R' and R denote the number of high-types out of $s+1$ and s SIVs respectively. Then the total payout by producers is

$$\begin{aligned}
\underline{t}_{s+1} &= \mathbb{E}\left(P'\Pi_{Q'}^{P'} + Q'\theta_L\widehat{\Delta}_{Q'}^{P'}\right) = \mathbb{E}\left(P'\Pi_{s+1-R'}^{P'}\right) + \mathbb{E}\left(Q'\theta_L\widehat{\Delta}_{Q'}^{P'}\right) \\
&= \mathbb{E}\left(\mathbb{E}\left(P'\Pi_{s+1-R'}^{P'} \mid R'\right)\right) + \mathbb{E}\left(\mathbb{E}\left(Q'\theta_L\widehat{\Delta}_{Q'}^{P'} \mid Q'\right)\right)
\end{aligned}$$

by the law of iterated expectations,

$$= \mathbb{E}\left(R'\theta_H\,\mathbb{E}\left(\Pi_{s+1-R'}^{P+1} \mid R'\right)\right) + \mathbb{E}\left(Q'\theta_L\mathbb{E}\left(\widehat{\Delta}_{Q'}^{P'} \mid Q'\right)\right)$$

by lemma 2.A.2 applied to $P' \mid R'$,

$$= (s+1)\gamma\theta_H\mathbb{E}\left(\Pi_{s-R}^{P+1}\right) + (s+1)(1-\gamma)\theta_L\mathbb{E}\left(\widehat{\Delta}_{Q+1}^{P}\right)$$

by lemma 2.A.2 applied to R' and Q',

$$= (s+1)(\gamma\theta_H + (1-\gamma)\theta_L)\mathbb{E}\left(\Pi_Q^{P+1}\right).$$

To show that (2.22) is increasing in γ, let us first rewrite it using the law of iterated expectations as

$$\begin{aligned}
\underline{t}_{s+1} &= (s+1)(\gamma\theta_H + (1-\gamma)\theta_L)\mathbb{E}\left(\Pi_Q^{P+1}\right) \\
&= (s+1)(\gamma\theta_H + (1-\gamma)\theta_L)\mathbb{E}\left(\mathbb{E}\left(\Pi_Q^{P+1} \mid Q\right)\right).
\end{aligned} \tag{2.41}$$

Since the marginal distribution of Q in (2.41) is $\text{Bin}(1-\gamma, s)$, by lemma 2.A.3, all we need to show is that $\mathbb{E}\left(\Pi_Q^{P+1} \mid Q\right)$ is monotonically decreasing in Q. To show this, let P_q and P_{q+1} represent $P \mid Q$ when $Q = q$ and $Q = q+1$, respectively. Then,

$$\Pi_q^{P_q+1} = \frac{1-\delta}{(1-\delta)(P_q+1)+\delta q} \tag{2.42}$$

$$= \frac{1-\delta}{(1-\delta)(P_{q+1}+1)+\delta q + (1-\delta)P} \quad \text{and} \tag{2.43}$$

$$\Pi_q^{P_{q+1}+1} = \frac{1-\delta}{(1-\delta)(P_{q+1}+1)+\delta(q+1)} \tag{2.44}$$

$$= \frac{1-\delta}{(1-\delta)(P_{q+1}+1)+\delta q + \delta} \tag{2.45}$$

where P is a Bernoulli random variable with parameter θ_H, since when there is one fewer defender there can be one more predator as long as he does not suffer a mutiny. Since $\delta > 1/2$, we see that (2.43) is always stricly greater than (2.45) so we have $\mathbb{E}\left(\Pi_Q^{P+1} \mid Q = q\right) > \mathbb{E}\left(\Pi_Q^{P+1} \mid Q = q+1\right)$, so that indeed $\mathbb{E}\left(\Pi_Q^{P+1} \mid Q\right)$ is decreasing in Q and $\square$

Appendix 2.B Tables

Table 2.1: OLS, not interacting with level of development

	1	2	3	4	5
Armed Forces	−0.215	0.064	0.030	0.203	0.210
	(0.25)	(0.40)	(0.41)	(0.42)	(0.39)
Income	2.921***	4.136***	4.039***	3.900***	3.860***
	(0.69)	(0.75)	(0.73)	(0.78)	(0.74)
Govt Expenditure		−0.796**	−0.936**	−0.932**	−1.050***
		(0.38)	(0.40)	(0.40)	(0.39)
Military Expenditure			0.360	0.280	0.244
			(0.27)	(0.27)	(0.26)
Population				−1.962	−2.275
				(1.74)	(1.74)
Internal Conflict					0.231**
					(0.09)
External Conflict					0.095
					(0.14)

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors reported in parentheses. All specifications with country and year fixed effects. Dependent variable: risk of expropriation.

Table 2.2: OLS, interacting with level of development

	1	2	3	4	5
Armed Forces	6.898***	6.774***	6.689***	6.977***	6.540***
	(1.19)	(1.17)	(1.17)	(1.23)	(1.18)
Income	5.197***	6.090***	6.013***	6.217***	6.062***
	(0.71)	(0.77)	(0.76)	(0.91)	(0.90)
Armed Forces × Income	−0.714***	−0.683***	−0.676***	−0.717***	−0.670***
	(0.12)	(0.12)	(0.12)	(0.14)	(0.13)
Govt Expenditure		−0.739**	−0.848**	−0.846**	−0.945**
		(0.37)	(0.39)	(0.38)	(0.37)
Military Expenditure			0.262	0.307	0.273
			(0.26)	(0.27)	(0.27)
Population				1.229	0.791
				(1.90)	(1.89)
Internal Conflict					0.196**
					(0.08)
External Conflict					0.039
					(0.13)

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors reported in parentheses. All specifications with country and year fixed effects. Dependent variable: risk of expropriation.

Table 2.3: OLS, interacting with OECD indicator

	1	2	3	4	5
Armed Forces	0.085	0.852**	0.820**	0.864**	0.852**
	(0.36)	(0.37)	(0.38)	(0.39)	(0.36)
Income	2.671***	4.095***	4.045***	3.994***	3.947***
	(0.68)	(0.71)	(0.70)	(0.75)	(0.71)
Armed Forces × OECD	−2.441***	−3.036***	−2.982***	−2.912***	−2.824***
	(0.63)	(0.63)	(0.63)	(0.65)	(0.64)
Govt Expenditure		−0.814**	−0.915**	−0.914**	−1.027***
		(0.37)	(0.39)	(0.39)	(0.37)
Military Expenditure			0.240	0.213	0.182
			(0.25)	(0.25)	(0.25)
Population				−0.720	−1.058
				(1.78)	(1.77)
Internal Conflict					0.216**
					(0.09)
External Conflict					0.099
					(0.13)

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors reported in parentheses. All specifications with country and year fixed effects. Dependent variable: risk of expropriation.

Table 2.4: OLS, sample restricted to 1995–2001

	1	2	3	4	5
Armed Forces	6.122***	4.840**	4.344*	5.948**	5.451**
	(2.26)	(2.35)	(2.37)	(2.65)	(2.50)
Income	6.069***	6.955***	6.574***	7.679***	7.255***
	(1.07)	(1.09)	(1.07)	(1.30)	(1.27)
Armed Forces × Income	−0.578**	−0.467*	−0.434*	−0.633**	−0.572**
	(0.23)	(0.24)	(0.24)	(0.29)	(0.27)
Govt Expenditure		−0.980**	−1.215***	−1.275***	−1.436***
		(0.44)	(0.44)	(0.44)	(0.47)
Military Expenditure			0.502*	0.582**	0.498*
			(0.28)	(0.29)	(0.28)
Population				5.161*	4.840*
				(2.90)	(2.79)
Internal Conflict					0.271**
					(0.11)
External Conflict					0.088
					(0.17)

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors reported in parentheses. All specifications with country and year fixed effects. Dependent variable: risk of expropriation.

Table 2.5: Instrumental variables

	1	2	3	4	5
Armed Forces	6.120***	6.777***	6.618***	6.764***	6.397***
	(1.48)	(1.33)	(1.34)	(1.41)	(1.44)
Income	4.765***	6.068***	6.055***	5.974***	5.842***
	(0.63)	(0.69)	(0.71)	(0.73)	(0.75)
Armed Forces × Income	−0.654***	−0.687***	−0.672***	−0.702***	−0.672***
	(0.14)	(0.13)	(0.13)	(0.14)	(0.14)
Govt Expenditure		−0.735***	−0.885***	−0.824***	−0.887***
		(0.27)	(0.28)	(0.28)	(0.28)
Military Expenditure			0.245	0.309	0.267
			(0.20)	(0.21)	(0.21)
Population				1.207	0.830
				(1.21)	(1.23)
Internal Conflict					0.174***
					(0.06)
External Conflict					0.024
					(0.08)

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors reported in parentheses. All specifications with country and year fixed effects. Dependent variable: risk of expropriation. The regressors are instrumented by their lags.

Table 2.6: Instrumental variables with OECD indicator

	1	2	3	4	5
Armed Forces	0.371	1.250***	1.246***	1.239***	1.124**
	(0.39)	(0.42)	(0.43)	(0.44)	(0.44)
Income	2.090***	3.811***	3.906***	3.798***	3.691***
	(0.50)	(0.58)	(0.59)	(0.58)	(0.59)
Armed Forces × OECD	−4.694***	−4.957***	−4.891***	−5.053***	−4.933***
	(0.98)	(0.92)	(0.94)	(1.00)	(1.00)
Govt Expenditure		−0.694**	−0.810***	−0.775***	−0.842***
		(0.28)	(0.28)	(0.28)	(0.28)
Military Expenditure			0.127	0.150	0.119
			(0.21)	(0.21)	(0.21)
Population				0.474	0.119
				(1.17)	(1.17)
Internal Conflict					0.174***
					(0.06)
External Conflict					0.081
					(0.08)

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors reported in parentheses. All specifications with country and year fixed effects. Dependent variable: risk of expropriation. The regressors are instrumented by their lags.

Table 2.7: OLS, lagged variables

	1	2	3	4	5
Armed Forces	4.482***	4.899***	4.836***	4.827***	4.782***
	(1.18)	(1.18)	(1.21)	(1.22)	(1.23)
Lagged Income	4.393***	6.152***	6.064***	5.826***	5.916***
	(0.94)	(1.03)	(1.07)	(1.15)	(1.15)
Armed Forces × Lagged Income	−0.629***	−0.610***	−0.605***	−0.588***	−0.578***
	(0.16)	(0.15)	(0.16)	(0.16)	(0.16)
Lagged Govt Expenditure		−1.209***	−1.247***	−1.226***	−1.283***
		(0.44)	(0.44)	(0.45)	(0.45)
Lagged Military Expenditure			0.144	0.095	0.070
			(0.38)	(0.38)	(0.39)
Lagged Population				−1.626	−1.782
				(1.95)	(1.97)
Lagged Internal Conflict					0.026
					(0.08)
Lagged External Conflict					−0.006
					(0.13)

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors reported in parentheses. All specifications with country and year fixed effects. Dependent variable: risk of expropriation.

Table 2.8: OLS, lagged variables with OECD indicator

	1	2	3	4	5
Armed Forces	0.148	1.158***	1.129**	1.172**	1.215***
	(0.33)	(0.43)	(0.45)	(0.46)	(0.46)
Lag Income	1.987***	4.051***	3.986***	3.881***	3.978***
	(0.75)	(0.82)	(0.84)	(0.90)	(0.88)
Armed Force × OECD	−2.258***	−3.040***	−3.001***	−2.903***	−2.906***
	(0.61)	(0.66)	(0.67)	(0.68)	(0.69)
Lagged Govt Expenditure		−1.232***	−1.271***	−1.258***	−1.332***
		(0.44)	(0.45)	(0.46)	(0.45)
Lagged Military Expenditure			0.136	0.109	0.089
			(0.35)	(0.36)	(0.36)
Lagged Population				−0.951	−1.081
				(2.01)	(2.03)
Lagged Internal Conflict					0.031
					(0.08)
Lagged External Conflict					0.029
					(0.13)

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors reported in parentheses. All specifications with country and year fixed effects. Dependent variable: risk of expropriation.

Chapter 3

Two-player rent-seeking contests with private values

3.1 Introduction

Following Tullock (1980)'s seminal contribution in the context of rent-seeking, there is now a large literature that models conflict in settings as varied as warfare, litigation and political competition, to name a few[1], as a *contest*, viz., a game where players expend effort or some other costly resources to increase the probability of winning a valuable prize. By and large, this literature has analysed conflict using contests under the assumption of complete information, i.e., there is common knowledge about how much contestants value the prize, or equivalently for the most commonly studied case of risk-neutral agents, how costly effort is. But the assumption of complete information implies that all parties agree on their respective probabilities of winning the prize and should therefore be able to come to a Coasian solution that avoids costly conflict altogether. For example, in the context of intellectual property rights litigation, if there is agreement on the value of a patent, say, all parties should be willing to agree to a settlement rather than engage in costly litigation in court.

Indeed, the conventional wisdom, articulated by, for example, Fearon (1995) and Wärneryd (2010, §5), is that it is precisely asymmetric information that generates conflict. To continue the above example, litigation might arise because parties possess information about the value of the patent that cannot be credibly revealed, because each side faces an incentive to exaggerate the value of the patent to them.

The natural way to model asymmetric information is the independent private values (IPV) formulation, where each risk-neutral player possesses a privately-known type describing his valuation of the prize, or equivalently, his marginal cost of effort and where these types are drawn from commonly-known independent distributions. The IPV model has been extensively studied in the case of perfectly discriminating contests[2] such as (first-price) all-pay auctions, but it has only recently begun receiving attention in the context of imperfectly discriminating

[1] For a survey, see Garfinkel and Skaperdas (2007a).

[2] Using the terminology of Hillman and Riley (1989), a perfectly discriminating contest is one where the player who exerts the highest effort wins the prize for sure. Imperfectly discriminating contests are ones where even expending the greatest effort does not guarantee winning the prize.

contests such as the commonly used Tullock or lottery contest[3].

The literature on the IPV model with Tullock contests of the ratio or power form[4] began with the study of binary distributions of type by Hurley and Shogren (1998), who only carried out numerical computations, and Malueg and Yates (2004), who considered a particular two-valued distribution that enables an analytical solution. Since then, the only contributions have been by Ryvkin (2010), who extends Fey (2008) to show existence of a pure-strategy equilibrium, and by Prada-Sarmiento (2010) and Wasser (2011) who compare rent-dissipation under complete and incomplete information. Apart from these general results, specific solutions to the equilibrium of Tullock contests have only been provided in purely numerical terms. The only exception is Ewerhart (2010), who provides the only known analytical solution arising from a particular distribution of types with bounded support.

The reason for the paucity of analytical solutions is that solving for the equilibrium strategies starting from a given distribution of types, what we shall call *forward solution*, poses formidable mathematical challenges as it involves solving a difficult integral equation. This paper proposes an alternative approach, which we shall call *reverse solution*, where we solve for the distribution of types that gives rise to a given distribution of efforts. The appeal of this approach is not solely mathematical, but stems also from the same considerations that make the revelead preference a valid approach to modelling consumer theory. To the extent that effort or resource expenditure in a contest is more readily observable than the value placed on a prize or the subjective cost of effort, it is reasonable to ask what must the latter be like in order for the former to satisfy certain properties.

Using the reverse solution method, this paper constructs a class of symmetric equilibria for the two-player case, where individual efforts are log-logistically distributed, deriving the distribution of types that gives rise to efforts of this kind. The log-logistic solution is of special interest as it is approximately equivalent, in a sense to be made precise later, to the equilibrium of a complete information contest. Also of note is the fact that these are the first examples in the literature of solutions where the distribution of efforts or types are not constrained to a bounded interval.

This reverse solution method also makes it possible to show that the equilibrium strategy, i.e., the mapping from type to effort, cannot take an affine form. This rules out two possible avenues for finding a forward solution, viz., assuming that the distribution of effort is of the same parametric family as that of types and plugging in an affine function into the integral equation given by first-order condition for the equilibrium.

The paper is structured as follows: section 3.2 lays out the standard two-player Tullock contest of the ratio form, describing the aforementioned two solution approaches, and proves the impossibility of affine strategies; section 3.3 constructs a distribution of types that give rise to log-logistically distributed efforts and carries out comparative statics with respect to the parameters of the distribution; section 3.4 discusses the results and provides concluding remarks. Proofs and results that are not essential to the flow of the paper are relegated to

[3]A Tullock or lottery contest is one where expending effort is akin to buying lottery tickets since the probability of winning is equal to the ratio of a player's effort to the total effort of all players.

[4]As opposed to those of the exponential form, about which even less is known in the case of asymmetric information.

appendix 3.A and appendix 3.B.

3.2 Model and general results

In a two-player Tullock or rent-seeking contest, risk-neutral players indexed by $i = 1, 2$ compete for a prize by expending effort x_i with constant marginal cost normalised[5] to 1. In the IPV model, the value that player i assigns to the prize is represented by a non-negative random variable V_i characterised by cumulative distribution function F_{V_i}, with V_i and V_j being independent. Each player i learns the realisation v_i of his value V_i and chooses simultaneously the level of effort x_i. We will consider contests of the ratio form, where the probability that player i wins the prize when he exerts effort x_i and player j exerts effort x_j is given by the contest success function

$$\pi_i(x_i, x_j) = \frac{x_i}{x_i + x_j}. \tag{3.1}$$

A pure strategy ξ_i of player i is a function that specifies the effort level $\xi_i(v_i)$ exerted when the private value is equal to v_i. The interim expected payoff of player i given that his valuation is v_i when exerts effort x_i and when player j adopts a strategy ξ_j is then given by

$$\begin{aligned} U_i(x_i; \xi_j, v_i) &= \mathbb{E}[v_i \pi_i(x_i, \xi_j(V_j)) - x_i] && (3.2)\\ &= v_i \mathbb{E}[\pi_i(x_i, \xi_j(V_j))] - x_i && (3.3)\\ &= v_i \int_0^\infty \frac{x_i}{x_i + \xi_j(v)}\, dF_{V_j}(v) - x_i. && (3.4) \end{aligned}$$

Differentiating under the integral, we obtain the derivative

$$U_i'(x_i; \xi_j, v_i) = v_i \int_0^\infty \frac{\xi_j(v)}{(x_i + \xi_j(v))^2}\, dF_{V_j}(v) - 1 \tag{3.5}$$

so that player i's best response $\hat{\xi}_i(v_i; \xi_j)$ to ξ_j satisfies the first-order condition

$$U_i'(\hat{\xi}_i(v_i; \xi_j); \xi_j, v_i) = v_i \int_0^\infty \frac{\xi_j(v)}{(\hat{\xi}_i(v_i; \xi_j) + \xi_j(v))^2}\, dF_{V_j}(v) - 1 = 0 \quad \text{for all } v_i > 0. \tag{3.6}$$

Therefore, a symmetric equilibrium where the players's values V_1 and V_2 are identically distributed and they use the same equilibrium strategy $\xi^*(v)$ must satisfy the first-order condition

$$\int_0^\infty \frac{\xi^*(u)}{(\xi^*(v) + \xi^*(u))^2}\, dF_V(u) = \frac{1}{v} \quad \text{for all } v > 0. \tag{3.7}$$

The *forward solution* for a symmetric pure strategy equilibrium therefore entails the formidable problem of solving for ξ^* in the integral equation (3.7), which is very poorly understood.

To avoid such difficulties we can rewrite the problem by writing $X_j = \xi_j(V_j)$ for the random variable that gives the effort level of player j, denoting the cdf of X_j by F_{X_j}, so that

[5] In the equivalent private cost formulation the value of the prize is normalised to 1 and the marginal cost of effort is private to each player.

player i's interim expected probability of winning the prize when exerting effort x_i as

$$P_{X_j}(x_i) = \mathbb{E}[\pi_i(x_i, X_j)] = \int_0^\infty \frac{x_i}{x_i + x}\, dF_{X_j}(x) \tag{3.8}$$

and the marginal probability of winning is given by its derivative

$$p_{X_j}(x_i) := P'_{X_j}(x_i) = \int_0^\infty \frac{x}{(x_i + x)^2}\, dF_{X_j}(x)\,. \tag{3.9}$$

The interim expected payoff is then

$$u_{X_j}(x_i; v_i) = v_i P_{X_j}(x_i) - x_i = v_i \int_0^\infty \frac{x_i}{x_i + x}\, dF_{X_j}(x) - x_i \tag{3.10}$$

and its derivative is

$$u'_{X_j}(x_i; v_i) = v_i p_{X_i}(x_i) - 1 = v_i \int_0^\infty \frac{x}{(x_i + x)^2}\, dF_{X_j}(x) - 1\,, \tag{3.11}$$

so that the best response $\hat{x}(v_i; X_j)$ to the opponent's effort X_j satisfies the first-order condition

$$u'_{X_j}(\hat{x}(v_i; X_j); v_i) = v_i p_{X_j}(\hat{x}(v_i; X_j)) - 1 = 0 \quad \text{for all } v_i > 0. \tag{3.12}$$

In a symmetric equilibrium where both players adopt the same strategy x^*, the equilibrium efforts $X_i^* = x^*(V_i)$ and $X_j^* = x^*(V_j)$ are independent and identically distributed random variables with the same cdf F_X. The equilibrium strategy x^* is then given by

$$x^*(v) = \hat{x}(v; X_i^*) = \hat{x}(v; X_j^*) = p_X^{-1}(1/v) \quad \text{where} \tag{3.13}$$

$$p_X(x) = \int_0^\infty \frac{t}{(x + t)^2}\, dF_X(t)\,, \tag{3.14}$$

whereas its inverse $x^{*-1}(x)$, giving the value corresponding to equilibrium effort x, is given by

$$1/v = p_X(x^*(v)) \qquad \Longleftrightarrow \qquad x^{*-1}(x) = 1/p_X(x)\,. \tag{3.15}$$

Then, letting V stand for either V_1 or V_2 and X for either X_1^* or X_2^*, the distribution of V is given by the cdf

$$F_V(v) = \Pr(V \le v) = \Pr\big(X \le x^*(v)\big) = \Pr\big(X_i \le p_X^{-1}(1/v)\big) = F_X\big(p_X^{-1}(1/v)\big)\,, \tag{3.16}$$

or alternatively, by the quantile function

$$Q_V(p) = F_V^{-1}(p) = \frac{1}{p_X\big(F_X^{-1}(p)\big)} = \frac{1}{p_X\big(Q_X(p)\big)}\,, \tag{3.17}$$

where Q_X is the quantile function of X, which also follows from the fact that

$$V = x^{*-1}(X) = \frac{1}{p_X(X)} . \tag{3.18}$$

This derivation leads to the following sufficient condition for an equilibrium.

Proposition 3.1 Let the random variables representing the values V_i of players $i = 1, 2$ be independent and identically distributed with cdf F_V, and suppose there exists a continuous random variable X with cdf F_X and pdf f_X such that $F_V(v) = F_X(p_X^{-1}(1/v))$ for all $v \in$ supp f_V, where

$$p_X(x) = \int_0^\infty \frac{t}{(x+t)^2} f_X(t)\, dt . \tag{3.19}$$

Then there exists a symmetric Bayesian equilibrium where effort is given by

$$x^*(v) = p_X^{-1}(1/v) \tag{3.20}$$

for all $v \in$ supp f_V, with F_X being the equilibrium distribution of effort.

Proof. Firstly, note that for p_X defined as in eq. (3.19),

$$p'_X(x) = \int_0^\infty \frac{-2t}{(x+t)^3} f_X(t)\, dt < 0 \quad \text{for all } x \geq 0, \tag{3.21}$$

holds for any pdf f_X so p_X is strictly monotonically decreasing. Therefore p_X has an inverse p_X^{-1} and the strategy x^* in eq. (3.20) is well-defined. Suppose now that player j is adopting the strategy x^*, so that player i's payoff when he values the prize at v and exerts effort x is

$$U(x; v) = v \int_0^\infty \frac{x}{x + x^*(u)}\, dF_V(u) - x \tag{3.22}$$

with marginal payoff

$$U'(x; v) = v \int_0^\infty \frac{x^*(u)}{(x + x^*(u))^2}\, dF_V(u) - 1 . \tag{3.23}$$

If we use the substitution $t = x^*(u) = p_X^{-1}(1/u)$, we can rewrite this as

$$U'(x; v) = v \int_0^\infty \frac{t}{(x+t)^2}\, dF_V(1/p_X(t)) - 1 \tag{3.24}$$

$$= v \int_0^\infty \frac{t}{(x+t)^2} f_X(t)\, dt - 1 \tag{3.25}$$

since by construction, $F_V(u) = F_X(p_X^{-1}(1/u))$, so that $U'(x; v) = v p_X(x) - 1$. Furthermore, by ineq. (3.21), $U''(x; v) = v p'_X(x) < 0$ for all $x \geq 0$, so that $U(x; v)$ is strictly concave in x and there is a unique x that maximises $U(x; v)$. Therefore, $x^*(v)$ is an equilibrium strategy if

$$U'(x^*(v); v) = v p_X(x^*(v)) - 1 = 0 \iff x^*(v) = p_X^{-1}(1/v) \tag{3.26}$$

holds for all $v \in \operatorname{supp} V$, which is true by construction. □

This proposition therefore gives us the a reverse method for finding a symmetric equilibrium:

Reverse solution method:

1. Choose a cdf F_X for the equilibrium level of effort that allows one to compute p_X in eq. (3.19).
2. If p_X has a closed-form inverse p_X^{-1}, find the equilibrium effort x^* using eq. (3.20) and the cdf of V using eq. (3.16).
3. Otherwise, find the inverse equilibrium effort x^{*-1} from eq. (3.15) and quantile function of V using eq. (3.17), as long as the quantile function Q_X has a closed-form expression.

To find a closed-form solution, therefore one needs to be able to evaluate the integral in eq. (3.19) analytically. This is obviously a much more mathematically tractable problem than the forward solution, as it involves plain integration rather than solving integral equations, and the examples in section 3.3 attest to this.

When putting this method in practice, it will often turn out to be the case that p_X cannot be inverted analytically. Although this means that we will be able to obtain a closed-form expression only for the inverse equilibrium effort x^{*-1}, this will still enable us to carry out comparative statics with respect to the parameters of the model. To do this, suppose that p_X and x^* depend on some parameter θ, then we can rewrite eq. (3.15) as

$$1/v = p_X(x^*(v;\theta);\theta) \tag{3.27}$$

and differentiate it with respect to θ, obtaining

$$\partial_\theta p_X(x^*(v;\theta);\theta) + \partial_x p_X(x^*(v;\theta);\theta)\, \partial_x x^*(v;\theta) = 0 \tag{3.28}$$

$$\iff \partial_\theta x^*(v;\theta) = -\frac{\partial_\theta p_X(x^*(v;\theta);\theta)}{\partial_x p_X(x^*(v;\theta);\theta)} \tag{3.29}$$

where ∂_x and ∂_θ represent partial derivatives with respect to the first and second arguments of x^* and p_X. Since p_X is monotonically decreasing by ineq. (3.21), the sign of the change in the equilibrium strategy due to a change in a parameter will be the same as that of the derivative of p_X with respect to that parameter.

In order to obtain an expression for p_X to be able to analyse the equilibrium, the choice of distribution for X is absolutely crucial. To aid in the choice of a suitable F_X in the above method, we will show that the probability P_X of winning satisfies certain distributional properties, but first we describe a distribution that will play a special role in the sequel.

Definition 3.1 A random variable X has log-logistic distribution[6] with scale parameter α and

[6]This distribution is also known as the Fisk distribution in economics. For more details see Kleiber and Kotz (2003, Chapter 6).

shape parameter β, written $X \sim \mathrm{LL}(\alpha, \beta)$, if it has cdf

$$F_X(x) = \frac{1}{1 + (x/\alpha)^{-\beta}} = \frac{x^\beta}{\alpha^\beta + x^\beta}. \tag{3.30}$$

The median is α and the mean is finite iff $\beta > 1$, in which case it is given by

$$\mathbb{E}(X) = \frac{\pi/\beta}{\sin(\pi/\beta)}. \tag{3.31}$$

The quantile function, or inverse cdf, is

$$Q_X(p) = \alpha \left(\frac{p}{1-p}\right)^{\frac{1}{\beta}}. \tag{3.32}$$

A standard log-logistic random variable is one with scale and shape parameters both equal to 1.

We can now state the following property of P_X.

Lemma 3.1 The probability $P_X(x)$ of winning the prize when exerting effort x and when the opponent's effort is given by the random variable X is equal to $\Pr(LX \leq x)$, where L is a standard log-logistic random variable independent of X.

PROOF. Let L be log-logistic with scale and shape parameters 1 so that its cdf is $F_L(\ell) = \frac{\ell}{1+\ell}$. Then then the cdf of LX is

$$\Pr(LX \leq x) = \Pr(L \leq x/X) = \mathbb{E}[\Pr(L \leq x/X \mid X)] = \mathbb{E}[F_L(x/X)] \tag{3.33}$$

$$= \int_0^\infty F_L(x/t)\, dF_X(t) = \int_0^\infty \frac{x/t}{1 + x/t}\, dF_X(t) = \int_0^\infty \frac{x}{x+t}\, dF_X(t), \tag{3.34}$$

which coincides with the definition of P_X in eq. (3.8). □

An immediate consequence of lemma 3.1 is that the marginal probability p_X that links the distribution of types and the equilibrium efforts is given by the pdf of LX.

Rather than working directly with the product LX and its corresponding pdf p_X, we will find it convenient to work with the sum $\log LX = \log L + \log X$ and its pdf ψ_X. Using the substitutions $x = e^\xi$, $t = e^\tau$, this leads to rewriting eq. (3.19) as

$$p_X(x) = \int_0^\infty \frac{t}{(x+t)^2} f_X(t)\, dt \tag{3.35}$$

$$\iff e^\xi p_X(e^\xi) = \int_{-\infty}^\infty \frac{e^{\xi-\tau}}{(1 + e^{\xi-\tau})^2} e^\tau f_X(e^\tau)\, d\tau \tag{3.36}$$

$$\iff \psi_X(\xi) = \int_{-\infty}^\infty \phi_L(\xi - \tau)\phi_X(\tau)\, d\tau \tag{3.37}$$

$$\iff \psi_X(\xi) = (\phi_L * \phi_X)(\xi) \tag{3.38}$$

where $*$ is the convolution operator and

$$\psi_X(\xi) = e^{\xi} p_X(e^{\xi}), \tag{3.39a}$$

$$\phi_X(\tau) = e^{\tau} f_X(e^{\tau}) \quad \text{is the pdf of } \log X \text{ and} \tag{3.39b}$$

$$\phi_L(z) = \frac{e^z}{(1+e^z)^2} \quad \text{is the standard logistic pdf.} \tag{3.39c}$$

This transformation can be useful for choosing a suitable distribution for X since it is often easier to work with the sum rather than the product of random variables. Furthermore, using the transformed functions in eq. (3.39), we can show the following result.

Proposition 3.2 Suppose that the distribution of values has a probability density function that is continuous on its support. Then the equilibrium effort cannot be an affine function of value.

PROOF. We will prove the result by contradiction. Suppose that $x^*(v) = a + bv$, with $b > 0$ since x^* is increasing. Since the distribution of V has a continuous pdf over its support, effort X will also have a continuous pdf, with a continuous pdf f over its support $[\underline{x}, \overline{x}]$, where we will write $\underline{x} = -\infty$ and $\overline{x} = \infty$ to mean that the support is unbounded below and above, respectively. Note then that the density ϕ of the log of effort is continuous and has support $[\underline{\xi}, \overline{\xi}]$, where $\underline{\xi} = \ln \underline{x}$ and $\overline{\xi} = \ln \overline{x}$.

By the first-order condition (3.20), $x^*(v) = p^{-1}(1/v)$ for all v in the support of V, so the function p must satisfy $p(x) = \frac{b}{x-a}$ for all $x \in [\underline{x}, \overline{x}]$, and the function $\psi(\xi) = e^{\xi} p(e^{\xi})$ must satisfy

$$\psi(\xi) = e^{\xi} p(e^{\xi}) = \frac{be^{\xi}}{e^{\xi} - a} = \frac{b}{1 - ae^{-\xi}} = 1 + \frac{b + ae^{-\xi}}{1 - ae^{-\xi}} \quad \text{for all } \xi \in [\underline{\xi}, \overline{\xi}]. \tag{3.40}$$

Suppose that X is unbounded above, i.e., $\overline{x}$ is infinite, then $\overline{\xi}$ is infinite so eq. (3.40) must be satisfied for all $\xi \in [\underline{\xi}, \infty)$. But $\psi(\xi) > 1$ for all $\xi > \ln a$, whereas by eq. (3.38), $\psi(\xi)$ is a density function over $[\underline{\xi}, \infty)$ which implies that $\psi(\xi) \to 0$ as $\xi \to \infty$. Hence $\overline{x}$ must be finite.

Suppose next that $a \leq 0$, so that $\frac{be^{\xi}}{e^{\xi}-a}$ is increasing everywhere if $a < 0$, or constant if $a = 0$. But, differentiating eq. (3.37) and evaluating at $\overline{\xi}$, we have

$$\psi'(\overline{\xi}) = \int_{\underline{\xi}}^{\overline{\xi}} \phi_L'(\overline{\xi} - \tau)\phi(\tau)\, d\tau < 0 \tag{3.41}$$

where $\phi_L(\ell) = \frac{e^{\ell}}{(1+e^{\ell})^2}$, since $\phi_L'(\ell) < 0$ for all $\ell > 0$ and since $\tau < \overline{\xi}$ inside the integral. Hence, $\psi' < 0$ for ξ sufficiently close to $\overline{\xi}$ by continuity of k' and ϕ. This means that ψ must be decreasing in the neighbourhood of $\overline{\xi}$ so that it cannot equal $\frac{be^{\xi}}{e^{\xi}-a}$ there, violating eq. (3.40). Hence it must be that $a > 0$. Since we must have $\underline{x} \geq a$ for $\frac{b}{x-a}$ to be non-negative, it follows that $\underline{x} > 0$.

The fact that $\underline{x} \geq a > 0$, implies that $\underline{\xi}$ is finite and that $\frac{be^{\xi}}{e^{\xi}-a}$ is decreasing for all $\xi > \underline{\xi}$.

On the other hand we have

$$\psi'(\underline{\xi}) = \int_{\underline{\xi}}^{\overline{\xi}} \phi'_L(\underline{\xi} - \tau)\phi(\tau)\, d\tau > 0 \tag{3.42}$$

since $\tau > \underline{\xi}$ inside the integral and $\phi'_L(\ell) > 0$ for $\ell < 0$, which means that $\psi' > 0$ for ξ sufficiently close to $\underline{\xi}$, by continuity of ϕ'_L and ϕ. This means that ψ must be increasing in ξ in the neighbourhood of $\underline{\xi}$ so that it cannot equal $\frac{be^{\xi}}{e^{\xi}-a}$ in that region, thus violating eq. (3.40). Hence for no f will ψ coincide with $\frac{be^{\xi}}{e^{\xi}-a}$ on its support. □

This implies that the distribution of effort will not in general be in the same family as the distribution of value. Furthermore, it rules out the most obvious way of attempting to solve for the equilibrium, namely by substituting the simplest possible equilibrium strategy.

3.3 Log-logistic model

3.3.1 Preliminaries

Suppose we take ϕ_X in eq. (3.37) to be a logistic pdf, so that ψ_X is the pdf of the sum of two logistic random variables. Since the logistic distribution is in many respects similar to the normal distribution, and since the sum of two normal random variables is itself normal[7], one might be tempted to think that the sum of two logistic random variables is also logistic, but this is not the case. On the other hand, the sum can be very closely approximated by a logistic random variable, as shown in appendix 3.A. Therefore, if we assume X to be log-logistically distributed then p_X will be very close to a log-logistic pdf, giving us the following.

Lemma 3.2 If $X \sim \mathrm{LL}(\alpha, \beta)$ then the probability and marginal probability of winning the prize are approximately given by

$$P_X(x) = \frac{x^{\gamma}}{\alpha^{\gamma} + x^{\gamma}} \tag{3.43}$$

$$p_X(x) = \frac{1}{x}\frac{\gamma\,(x/\alpha)^{\gamma}}{(1 + (x/\alpha)^{\gamma})^2}, \tag{3.44}$$

where $\gamma := \frac{1}{\sqrt{1+1/\beta^2}}$ is strictly less than 1 for all $\beta > 0$.

Proof. See proof 3.A.1 in appendix 3.A. □

Since the analysis in appendix 3.A indicates that the approximation is very precise, we will use the expressions in eq. (3.43) and eq. (3.44) as if they were exact for P_X and p_X, without necessarily qualifying them as being approximate.

[7]This suggests using the normal distribution as an approximation, so that p_X would be a log-normal pdf. But such an approximation would be a *qualitatively* poor one since for no value of its parameters is a log-normal pdf monotonically decreasing over its support, as required by ineq. (3.21). For similar reasons, although the contest success function $\pi_i(x_i, x_j) = \Phi(\log(x_i/x_j))$ would give us an exact expression for the marginal probability of winning the prize when effort is log-normally distributed, this would mean that the expected payoff is not concave in own effort, leading to problems when using the first-order condition.

Note that, since $\gamma < 1$ for all $\beta > 0$, the marginal probability of winning p_X is monotonically decreasing for all $x > 0$, which is consistent with ineq. (3.21). Also p_X becomes arbitrarily large as x tends to 0 and tends to 0 as $x \to \infty$. Since $V = 1/p_X(X)$, by eq. (3.18), this means that the distribution of values has support equal to $[0, \infty)$. In fact, using eq. (3.18) and eq. (3.44) we can write

$$V = \frac{1}{p_X(X)} = \frac{X}{\gamma}\frac{(1 + (X/\alpha)^\gamma)^2}{(X/\alpha)^\gamma} \tag{3.45}$$

$$= \frac{X}{\gamma}\left(2 + (X/\alpha)^{-\gamma} + (X/\alpha)^\gamma\right) \tag{3.46}$$

$$= \frac{1}{\gamma}(2X + \alpha^\gamma X^{1-\gamma} + \alpha^{-\gamma} X^{1+\gamma}) \,. \tag{3.47}$$

Before continuining our characterisation of V we will need the following result.

Lemma 3.3 If $X \sim \mathrm{LL}(\alpha, \beta)$ then $X^\gamma \sim \mathrm{LL}(\alpha^\gamma, \beta/\gamma)$.

PROOF. See proof 3.B.1 in appendix 3.B. □

Combining eq. (3.47), lemma 3.3 and eq. (3.32), we find that the quantile function Q_V of V is

$$Q_V(p) = \frac{1}{\gamma}\left(\alpha^\gamma \alpha^{1-\gamma}\left(\frac{p}{1-p}\right)^{\frac{1-\gamma}{\beta}} + \alpha^{-\gamma}\alpha^{1+\gamma}\left(\frac{p}{1-p}\right)^{\frac{1+\gamma}{\beta}} + 2\alpha\left(\frac{p}{1-p}\right)^{\frac{1}{\beta}}\right) \tag{3.48}$$

$$= \frac{\alpha}{\gamma}\left(\left(\frac{p}{1-p}\right)^{(1+\gamma)\tilde{\gamma}} + \left(\frac{p}{1-p}\right)^{(1-\gamma)\tilde{\gamma}} + 2\left(\frac{p}{1-p}\right)^{\tilde{\gamma}}\right) \,, \tag{3.49}$$

where $\tilde{\gamma} = 1/\beta = \sqrt{1-\gamma^2}/\gamma$.

3.3.2 Equilibrium

We can now define the distribution of values that will give the desired equilibrium.

Definition 3.2 A random variable V has VL distribution with scale parameter $\alpha > 0$ and shape parameter $\gamma \in (0, 1)$, written $V \sim \mathrm{VL}(\alpha, \gamma)$, iff

$$V = \frac{\alpha}{\gamma}\left(Y^{(1-\gamma)\tilde{\gamma}} + 2Y^{\tilde{\gamma}} + Y^{(1+\gamma)\tilde{\gamma}}\right) \,, \tag{3.50}$$

where $\tilde{\gamma} := \sqrt{1-\gamma^2}/\gamma$ and Y is a standard logistic random variable.

Lemma 3.4 If $V \sim \mathrm{VL}(\alpha, \gamma)$,
(a) the pdf f_V of V is unimodal;
(b) the median of V is $4\alpha/\gamma$;
(c) the expectation of V is

$$\mathbb{E}(V) = \frac{\alpha}{\gamma}\left(\frac{\pi(1-\gamma)\tilde{\gamma}}{\sin(\pi(1-\gamma)\tilde{\gamma})} + 2\frac{\pi\tilde{\gamma}}{\sin(\pi\tilde{\gamma})} + \frac{\pi(1+\gamma)\tilde{\gamma}}{\sin(\pi(1+\gamma)\tilde{\gamma})}\right) \,, \tag{3.51}$$

if $\gamma > \gamma^*$, where $\gamma^* \approx 0.883$ is the unique positive root of $\gamma^4 + 2\gamma^3 + \gamma^2 - 2\gamma - 1 = 0$.

Figure 3.1: Plot of density $f_V(v)$ of prize values

$\alpha = 1$ and $\gamma = 1/2, 2/3, 4/5, 9/10$ (dashed), $\gamma = 0.95, 0.975, 0.99$ (solid). Curves with modes further to the right correspond to higher γ.

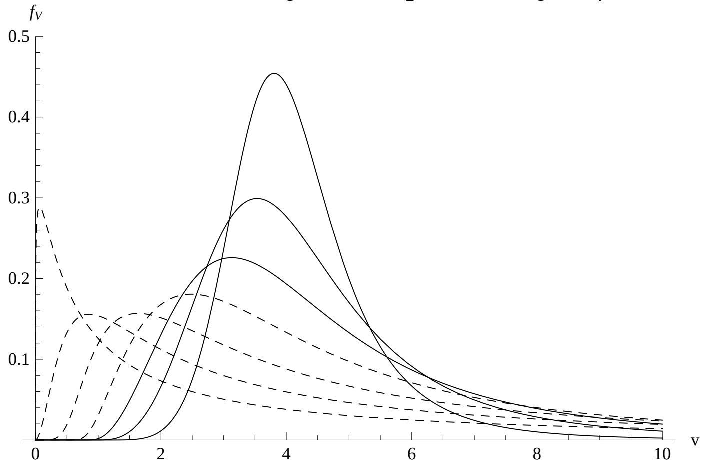

PROOF. See proof 3.B.4 in appendix 3.B. □

Having derived the distribution of values that leads to log-logistic efforts we can now formally describe the equilibrium.

Proposition 3.3 If the valuations V_1 and V_2 are i.i.d., with $V_i \sim \mathrm{VL}(\alpha, \gamma)$, $\alpha > 0$, $0 < \gamma < 1$, the inverse of the equilibrium strategy is approximately given by

$$x^{*-1}(x) = \frac{x}{\gamma}\left(\left(\frac{x}{\alpha}\right)^{-\gamma} + 2 + \left(\frac{x}{\alpha}\right)^{\gamma}\right), \tag{3.52}$$

with the equilibrium efforts X_1^* and X_2^* having distribution $X_i^* \sim \mathrm{LL}(\alpha, \beta)$, where $\beta = \gamma/\sqrt{1-\gamma^2}$.

PROOF. By construction, we can apply proposition 3.1 with p_X given by lemma 3.2. □

Unfortunately, it is not in general possible to invert x^{*-1} in eq. (3.52) to obtain a closed-form expression for $x^*(v)$, but for rational γ, x^* corresponds to a root of a polynomial so that it can be plotted with arbitrary precision[8], as is done in fig. 3.2.

3.3.3 Comparative statics

In order to analyse how the equilibrium varies as the distribution of values changes, we first need to understand the role that the parameters α and γ play. From the definition of V in eq. (3.50) it is clear that α is indeed a simple scale parameter, so that higher α leads to V taking proportionally larger values. In fact, since X and V share the same scale parameter α, any quantities that are homogeneous functions of X and V will be proportional to α. On the other hand, the role of γ is not so obvious.

[8] By the Abel-Ruffini theorem, there is no general algebraic solution to polynomials of degree greater than 4, but the numerical analysis of roots of polynomials is well-understood, to the extent that all major computer algebra systems offer reliable tools for manipulating them.

Figure 3.2: Plot of equilibrium effort $x^*(v)$

$\alpha = 1$ and $\gamma = 1/2, 2/3, 4/5, 9/10, 0.95, 0.99$. Curves with higher right endpoints correspond to higher γ. The dashed line represents the limit as $\gamma \to 1$, given by $\sqrt{v} - 1$.

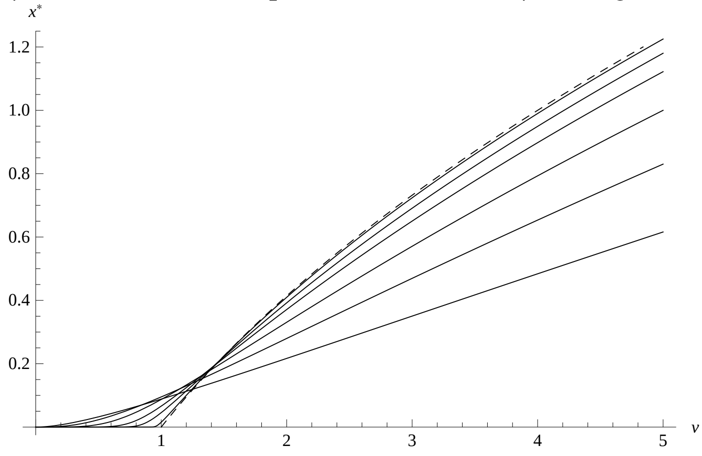

Inspecting fig. 3.1, greater values of γ seem to correspond to a distribution that is less disperse. In order to make this more precise we will use the Lorenz ordering of distributions, defined as follows.

Definition 3.3 If X_1 and X_2 are two random variables with finite expectations, we say that X_1 exhibits less inequality than X_2 in the Lorenz sense, denoted $X_1 \leq_L X_2$, if $L_{X_1}(p) \leq L_{X_2}(p)$ for all $p \in [0, 1]$, where L_{X_i} is the Lorenz curve of X_i defined by

$$L_{X_i}(p) = \frac{1}{\mathbb{E}(X_i)} \int_0^p Q_{X_i}(q)\, dq \tag{3.53}$$

and Q_{X_i} is the quantile function of X_i.

Since Atkinson (1970)'s seminal paper, the Lorenz ordering has become one of the standard criteria used to compare income distributions in terms of their inequality. Applied to the present context, saying that $X_1 \leq_L X_2$ means that the distribution of X_1 is less unequal or less disperse than that of X_2. The Lorenz ordering can only rank distributions whose Lorenz curves do not cross, but in the case of log-logistic distributions, the order turns out to be total.

Lemma 3.5 Suppose $X_1 \sim \mathrm{LL}(\alpha_1, \beta_1)$ and $X_2 \sim \mathrm{LL}(\alpha_2, \beta_2)$, with $\beta_1 > 1$ and $\beta_2 > 1$. Then $\beta_1 \geq \beta_2 \iff X_1 \leq_L X_2$. Also, if $X \sim \mathrm{LL}(\alpha, \beta)$, the Gini coefficient for X is $1/\beta$.

PROOF. See Kleiber and Kotz (2003, p. 224) and the concluding remarks in Wilfling (1996).□

This shows that log-logistic distributions are completely ordered in the Lorenz-sense by the shape parameter β and that $1/\beta$ is a direct measure of dispersion. Indeed, as $\beta \to \infty$, the distribution degenerates into a point mass at α. It also turns out the shape parameter γ has an equivalent role in the VL distribution in definition 3.2.

Lemma 3.6 Suppose $V_1 \sim \mathrm{VL}(\alpha_1, \gamma_1)$ and $V_2 \sim \mathrm{VL}(\alpha_2, \gamma_2)$, with both $\mathbb{E}(V_1)$ and $\mathbb{E}(V_2)$ finite. Then $\gamma_1 \geq \gamma_2 \iff V_1 \leq_L V_2$.

Figure 3.3: Plot of quantile function $Q_X(p)$ of effort

$\alpha = 1$ and $\gamma = 1/2, 2/3, 4/5, 9/10, 0.95, 0.99$. Higher values of γ correspond to curves that are lower for $p < 1/2$ and higher for $p > 1/2$.

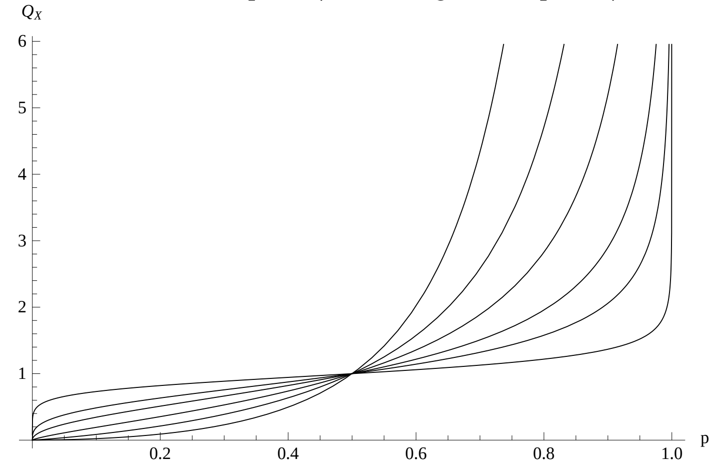

PROOF. See proof 3.B.5 in appendix 3.B. □

Hence, the family of VL distributions is also totally ordered in the Lorenz sense by γ, with higher γ indicating lower dispersion. We therefore have the following comparative statics.

Proposition 3.4 As α increases, the distributions of values and equilibrium efforts scale up in the same proportion. As γ increases, the distribution of effort become less disperse.

PROOF. The first statement is obvious because α is the same scale parameter of both distributions. For the second statement, by proposition 3.3, the shape parameter of the log-logistic equilibrium distribution of effort is $\beta = \gamma/\sqrt{1-\gamma^2}$, which is increasing in γ. Therefore, by lemma 3.5, the equilibrium distribution of effort becomes less disperse. □

Figure 3.3 plots the quantile functions of equilibrium effort as γ increases, showing how they become flatter and flatter, tending to 1 as $\gamma \to 1$.

Remark 3.1 As $\gamma \to 1$, $X \xrightarrow{p} V/4$.

PROOF. As $\gamma \to 1$, all the powers in eq. (3.50) tend to 0, so that V tends in probability to 4α, whereas $\beta \to \infty$ so that X tends in probability to α. □

This shows that as uncertainty about values disappears, the equilibrium strategy tends towards that under complete information[9], as we would expect.

We can also analyse how changes in α and γ affect the interim equilibrium efforts, i.e., effort once the value of the prize is realised. Although we cannot invert eq. (3.52) analytically to obtain x^*, we can still use eq. (3.29) to determine carry out comparative statics.

Proposition 3.5 For constant value v, a small increase in α, i.e., a small proportional increase in the distribution of values, leads to an increase in equilibrium effort if the value is above the median and a decrease if the value is below the median.

[9]See Wärneryd (2010) for a derivation of the complete information equilibrium.

PROOF. Using eq. (3.29) we know that the sign of $\partial_\alpha x^*(v;\alpha,\gamma)$ is the same as that of

$$\partial_\alpha p_X(x;\alpha,\gamma) = \frac{\partial}{\partial\alpha}\frac{1}{x}\frac{\gamma\,(x/\alpha)^\gamma}{(1+(x/\alpha)^\gamma)^2} = \gamma^2\frac{\alpha^{\gamma-1}x^{\gamma-1}}{(\alpha^\gamma+x^\gamma)^3}(x^\gamma-\alpha^\gamma) \tag{3.54}$$

which is positive if $x > \alpha$ and negative if $x < \alpha$. Since α is the median of the equilibrium distribution of effort, this means that $\partial_\alpha x^*$ is positive if the realised value is above the median and negative if it's below. □

This result shows that scaling up the distribution of values by a small amount increases effort in equilibrium only if a player has a realised value that is above the median, since that is when a scaling up will make it more likely that the opponent's value is close the player's thus increasing the marginal benefit of effort.

Proposition 3.6 For constant value v, a small increase in γ, i.e., a small decrease in the dispersion of the distribution of values, leads to a decrease in equilibrium effort if the value is in the top or bottom $\delta(\gamma)$ quantile and to an increase in effort otherwise, where $\delta(\gamma) \in (0, 1/2]$ is a decreasing function.

PROOF. See proof 3.B.6 in appendix 3.B. □

This states that when uncertainty about values decreases by a small amount, a contestant who has an interim value in some interval around the median increases his effort, since the probability that the opponent's value is close to his increases, thus increasing the marginal benefit of effort. Figure 3.4 depicts the conditions under which effort increases with a reduction in dispersion.[10]

Figure 3.4: Sign of $\partial_\gamma x^*$

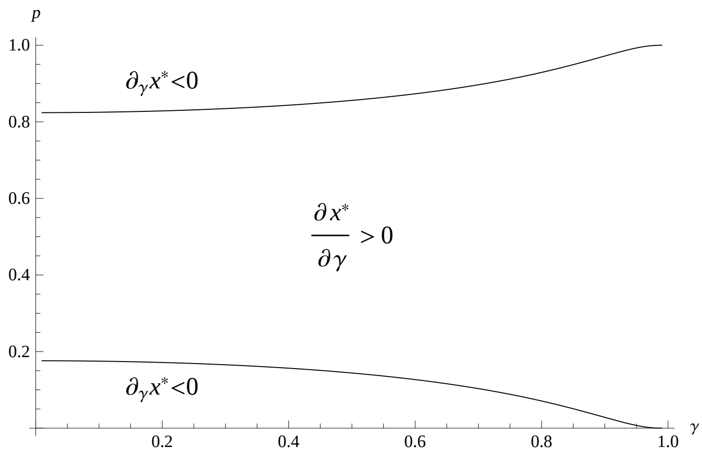

Proposition 3.7 For $\gamma > \gamma^*$, a mean-preserving spread in the distribution of values decreases expected equilibrium effort.

[10]Note that the δ quantiles of V are so large that they are not in the range of values in fig. 3.2, which does not show x^* decreasing for large v.

PROOF. See proof 3.B.7 in appendix 3.B. □

Figure 3.5: Normalised mean effort $\mathbb{E}X/\mathbb{E}V$

$\gamma \in (\gamma^*, 1)$

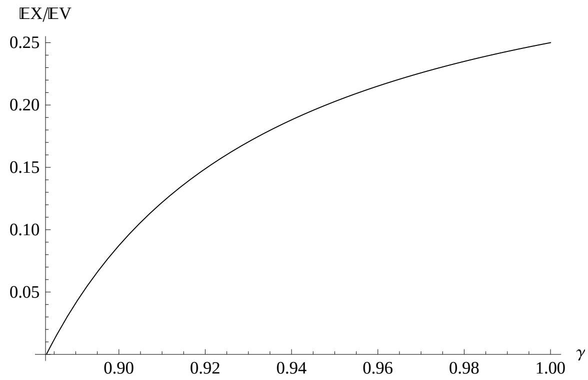

A mean-preserving spread in V consists of a fall in γ accompanied by a fall in α. As $\gamma \to \gamma^*$ from above, $\mathbb{E}V$ becomes arbitrarily large, so that to keep the mean constant, $\alpha \to 0$, which explains why $\mathbb{E}X \to 0$ as well in fig. 3.5.

An important quantity in the study of rent-seeking contests is the *willingness to waste*, defined as the ratio of the effort expenditure to the value of the prize, which measures the amount of rent that is dissipated by each player. For the present case of log-logistic effort we have the following result.

Proposition 3.8 The willingness to waste is a hump-shaped function of value, attaining its maximum of $\gamma/4$ when $v = 4\alpha/\gamma$. Its quantile function is given by

$$w_X(p) = \frac{\gamma\big((1-p)p\big)^{\sqrt{1-\gamma^2}}}{\big((1-p)^{\sqrt{1-\gamma^2}} + p^{\sqrt{1-\gamma^2}}\big)^2}, \tag{3.55}$$

which is increasing in γ for all p.

PROOF. See proof 3.B.8 in appendix 3.B. □

This result shows that the willingness to waste is highest when a contestant has the median value for the prize, since this is when he is most likely to be faced with an evenly matched opponent. As the realised value becomes more and more extreme, the contest becomes more and more lopsided, decreasing the incentive to exert effort.

Furthermore, as γ increases, uncertainty decreases and the distribution of the willingness to waste increases uniformly, in the sense that all its quantiles increase, as shown in fig. 3.6. This means that greater uncertainty reduces rent-dissipation by reducing the incentives to expend effort.

3.3.4 Correspondence between incomplete and complete information

To develop an intuition for this result it is useful to look at the expression in eq. (3.43) (reproduced below) for the probability of winning the prize when the opponent is following the

Figure 3.6: Willingness to waste

$\gamma = 1/2, 2/3, 4/5, 9/10, 0.95, 0.99$. Higher curves correspond to higher γ.

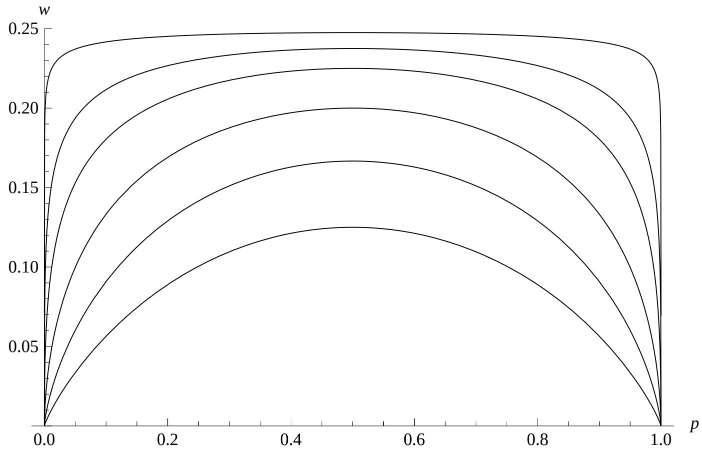

equilibrium strategy, which is

$$P_X(x) = \frac{x^\gamma}{x^\gamma + \alpha^\gamma}. \tag{3.43}$$

Consider now a *complete* information contest where the contest success function giving the probability that a player wins when he exerts effort x and his opponent exerts effort y is of the generalized ratio-form

$$\pi(x; y) = \frac{x^\gamma}{x^\gamma + y^\gamma}. \tag{3.56}$$

Then, it is clear that the probability of winning when faced with an opponent whose prize value is distributed according to definition 3.2 and who is following the equilibrium strategy is equal to the probability of winning in a generalized complete information contest when faced with an opponent who exerts effort α. Moreover, since α is the median level of equilibrium effort under incomplete information, the contestants' incentives are as if they were facing opponents of the median type. More precisely, this means that the equilibrium strategy $x^*(v)$ under incomplete information is identical to the *reaction function* under complete information but with the generalized contest success function in eq. (3.56).

This equivalence makes the effect of the parameter γ intuitively clear: as γ decreases, the probability of winning becomes less and less responsive to effort. Ceteris paribus, this decreases the incentive to expend effort so that in the limit as $\gamma \to 0$ effort expenditure tends to 0. Conversely, as $\gamma \to 1$, the probability of winning becomes more and more responsive to effort, which we can see also in the shape of the equilibrium strategy $x^*(v)$ as depicted in fig. 3.2. Indeed, as $\gamma \to 1$, $x^*(v) \to \sqrt{v} - \alpha$, which is the reaction function under complete information.

3.4 Discussion and conclusion

The log-logistic model in section 3.3 shows how the reverse solution method proposed in section 3.2 can be put to use to find an equilibrium of a two-player incomplete information contest in analytical, rather than purely numerical, form. The attractions of the log-logistic model thus solved are several. Firstly, the model exhibits a formal correspondence with a suitably generalized complete information contest, in the sense that each player can replace the probability of winning, which is an expectation over the opponent's effort, by the probability of winning against an opponent of the median type following the equilibrium strategy. Thi s shows that the equilibrium with log-logistic effort we have examined is not an arbitrary example chosen merely for its analytical tractability; rather, section 3.3.4 demonstrates a more fundamental connection with contests under complete information.

Secondly, even though only the inverse of the equilibrium strategy has a closed-form expression and not the equilibrium strategy itself, the quantiles of the distribution of values, of effort and of the willingness to waste have simple functional forms. In fact, the lack of a closed-form expression for the equilibrium strategy is a direct consequence of the fact that *even under complete information*, there is no closed-form expression for reaction function of each player when the power in the contest-success function is less than 1, and is not a complication introduced by incomplete information.

Thirdly, the comparative statics of the equilibrium are particularly simple and intuitive, to wit:

1. scaling the distribution of values up, scales the distribution of efforts up by the same proportion, keeping rent dissipation constant;
2. increasing the dispersion of the distribution of values, increases the dispersion of the distribution of efforts, where dispersion is measured according to the Lorenz criterion;
3. increasing uncertainty in the values decreases mean effort expenditure and rent dissipation.

These results contrast with the only other existing analytical solution to a Tullock contest under incomplete information, namely that of Ewerhart (2010)[11], where changes in the parameters of the distribution of values do not yield simple, unambiguous changes in the distribution of effort. Also of note, is the fact that this paper provides the first solution, numerical or analytical, where the distribution of values and that of effort has support over the whole of the positive real line, rather than on some bounded interval.

The major limitations of the approach in the present paper are that it is restricted to the case of two players only and to symmetric equilibria where the two players have the same distribution of values. Although symmetric distributions of values are standard, the literature has studied the general case of more than two players. Extending the paper's approach, and the log-logistic model to the general case of several players should provide an avenue for fruitful reseach.

[11] In fact, the solution obtained by Ewerhart (2010) can be derived using the method presented here by letting ϕ_X be uniform, so that X is log-uniformly distributed.

Appendix 3.A Accuracy of the logistic approximation

In this section, we study the accuracy of the approximation used in lemma 3.2.

Except for certain special cases, it is not possible to derive closed-form expressions for the cdf or the pdf of the sum of logistic random variables, so we will instead have to rely on characteristic functions[12], defined as follows.

Definition 3.A.1 The characteristic function $\varphi_X : \mathbb{R} \to \mathbb{C}$ of a random variable X is defined by $\varphi_X(t) = \mathbb{E}(e^{itX})$, and is equal to

$$\varphi_X(t) = \int_{-\infty}^{\infty} e^{itx} f_X(x)\, dx \tag{3.57}$$

if X has a pdf f_X.

The usefulness of characteristic functions in the present context lies in the fact that the characteristic function of the sum of two independent random variables is the product of the characteristic functions of the random variables, which follows immediately from the fact that

$$\mathbb{E}\big(e^{it(X+Y)}\big) = \mathbb{E}\big(e^{itX}e^{itY}\big) = \mathbb{E}\big(e^{itX}\big)\mathbb{E}\big(e^{itY}\big), \tag{3.58}$$

by independence of X and Y.

Before proceding further, let us define a parametrization of the logistic function and state some of its properties that we will need.

Definition 3.A.2 If a random variable X has logistic distribution with location parameter μ and scale parameter s, written $X \sim \mathrm{L}(\mu, s)$, its cdf and characteristic function are

$$F_X(x) = \frac{1}{1 + e^{-(x-\mu)/s}} \tag{3.59}$$

$$\text{and} \quad \varphi_X(t) = e^{it\mu}\frac{\pi s t}{\sinh(\pi s t)}, \tag{3.60}$$

respectively. Its mean is μ and its variance is $\frac{\pi^2}{3}s^2$.

If $\widetilde{L} = \log L \sim \mathrm{L}(0, 1)$ and $\widetilde{X} = \log X \sim \mathrm{L}(\mu, s)$ are independent, then $\mathbb{E}(\widetilde{L}+\widetilde{X}) = \mu$ and $\mathrm{Var}(\widetilde{L}+\widetilde{X}) = \frac{\pi^2}{3}(1+s^2)$, which suggests using $\widetilde{L+X} \sim \mathrm{L}\left(\mu, \sqrt{1+s^2}\right)$ as an approximation to $\widetilde{L} + \widetilde{X}$. The next result establishes the upper bound on the error in this approximation.

Lemma 3.A.1 Let $\widetilde{L} \sim \mathrm{L}(0, 1)$, $\widetilde{X} \sim \mathrm{L}(\mu, s)$ be independent and let $\widetilde{L+X} \sim \mathrm{L}\left(\mu, \sqrt{1+s^2}\right)$. Then the absolute difference Δ of the characteristic functions $\widetilde{\varphi}$ and φ of $\widetilde{L+X}$ and $\widetilde{L} + \widetilde{X}$, respectively, is

$$\Delta = |\widetilde{\varphi} - \varphi| \le 0.0355072 \tag{3.61}$$

for all $t \in \mathbb{R}$ and for all $s > 0$, where the maximum is attained when $s = 1$.

[12] See Billingsley (1995, p. 342) for more details.

Figure 3.7: Absolute error Δ of characteristic function approximation

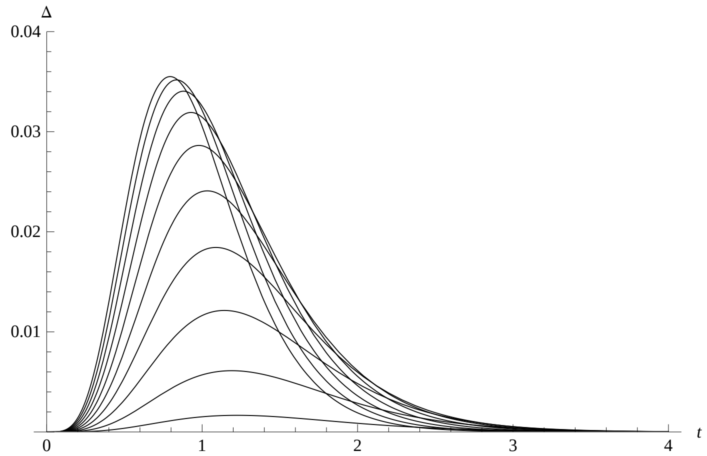

(a) $s = 1, 0.9, 0.8, 0.7, 0.6, 0.5, 0.4, 0.3, 0.2, 0.1$. Curves with higher peaks correspond to higher s.

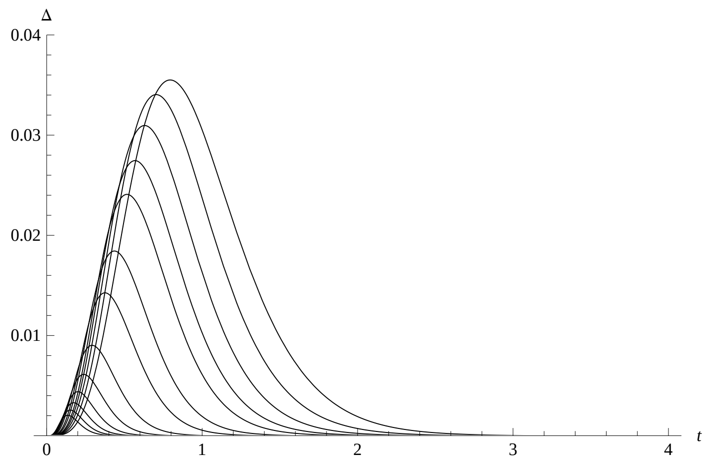

(b) $s = 1, 1.25, 1.5, 1.75, 2, 2.5, 3, 4, 5, 6, 7, 8, 9$. Curves with higher peaks correspond to lower s.

PROOF. See proof 3.B.2 in appendix 3.B. □

Furthermore, fig. 3.7 shows that the error falls rapidly from the upper bound both as t increases and as s moves away from 1.

This approximation gives us a formula for ψ_X, but we are really after is p_X. Having assumed that $\widetilde{X} = \log X$ is logistic, it follows that $X = e^{\widetilde{X}}$ is log-logistic and $LX = e^{\widetilde{L}+\widetilde{X}} \overset{d}{\approx} e^{\widetilde{L+X}}$ is approximately log-logistic. We can then establish lemma 3.2.

Proof 3.A.1 (Proof of lemma 3.2) Comparing definition 3.1 and definition 3.A.2, we see that if X is log-logistic with scale parameter α and shape parameter β, $\widetilde{X} = \log X$ has location parameter $\log\alpha$ and scale parameter $1/\beta$. Therefore, the logistic approximation $\widetilde{L+X}$ to $\widetilde{L} + \widetilde{X}$ has location parameter $\log\alpha$ and scale parameter $\sqrt{1+\beta^{-2}}$, so that the log-logistic random variable $e^{\widetilde{L+X}}$ has scale parameter α and shape parameter $1/\sqrt{1+\beta^{-2}}$. Using the approximation $\widetilde{L+X} \overset{d}{\approx} \widetilde{L} + \widetilde{X}$, we have $e^{\widetilde{L+X}} \overset{d}{\approx} e^{\widetilde{L}+\widetilde{X}} = LX$ so that by lemma 3.1, P_X and p_X are approximately the cdf and pdf of a log-logistic random variable with the given parameters. ■

We should check that the approximation remains valid under the transformation from ψ_X to p_X. Fortunately, it turns out that P_X and p_X can be computed exactly for the special case of $\alpha = \beta = 1$.

Lemma 3.A.2 If $X \sim \mathrm{LL}(1, 1)$ then

$$P_X(x) = \frac{x(x-1-\ln x)}{(x-1)^2} \tag{3.62}$$

$$\text{and} \quad p_X(x) = \frac{(1+x)\ln x - 2(x-1)}{(x-1)^3}. \tag{3.63}$$

PROOF. See proof 3.B.3 □

Recall from lemma 3.A.1 that the maximum error in the logistic approximation occurs when scale parameter of $\log X$ is 1. Since this corresponds to the shape parameter of X being 1, we can compare the exact expressions in lemma 3.A.2 with the approximate ones in lemma 3.2 for the worst case $\alpha = \beta = 1$. Visual inspection of the plot of the two in fig. 3.8 suggests that the approximation is rather accurate. Confirming this, panels (a) and (b) in fig. 3.9 show that for values of x greater than about 0.2, the absolute error Δ is less than about 0.015 in magnitude, declining very rapidly as x increases. On the other hand, since both the exact and approximate expressions for p_X become arbitrarily large as x tends to 0, the absolute error also becomes arbitrarily large. In this case, the relative error δ provides a better indicator of the accuracy of the approximation. Panels (c) and (d) show that for x larger than 0.002 this is at most 0.1.

Figure 3.8: Comparison of exact and approximate expressions for p_X

$\alpha = 1, \beta = 1$. Solid line: exact, dashed line: approximate

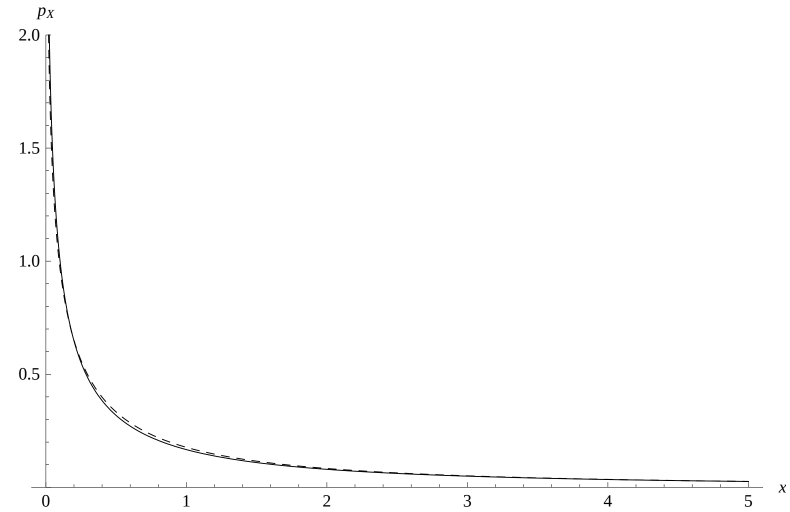

Figure 3.9: Approximation error for p_X

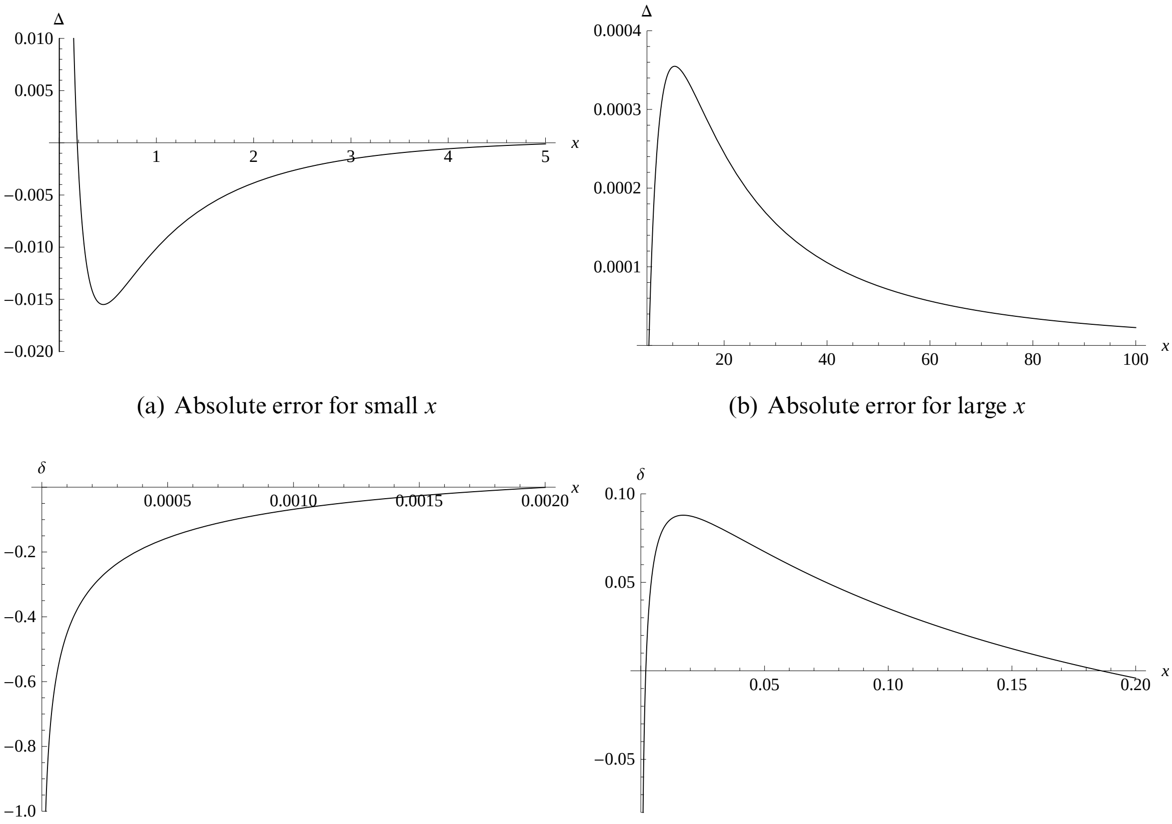

(a) Absolute error for small x

(b) Absolute error for large x

(c) Relative error for very small x

(d) Relative error for small x

Appendix 3.B Proofs

Proof 3.B.1 (Proof of lemma 3.3) If $X \sim \mathrm{LL}(\alpha, \beta)$ then $\log X \sim \mathrm{L}(\log\alpha, 1/\beta)$. Since $\log X^\gamma = \gamma \log X$ and

$$\Pr(\gamma \log X \leq x) = \Pr(\log X \leq x/\gamma) = \frac{1}{1 + e^{-(x/\gamma - \log\alpha)\beta}} \tag{3.64}$$

$$= \frac{1}{1 + e^{-(x - \gamma\log\alpha)\beta/\gamma}} = F_{\mathrm{L}(\log\alpha^\gamma, b/\gamma)}(x)\,, \tag{3.65}$$

it follows that $X^\gamma \sim \mathrm{LL}(\alpha^\gamma, \beta/\gamma)$. ■

Proof 3.B.2 (Proof of lemma 3.A.1) The characteristic functions of $\widetilde{L+X}$ and $\widetilde{L} + \widetilde{X}$ are given by

$$\widetilde{\varphi}(t) = e^{it\mu} \frac{\pi t\sqrt{1+s^2}}{\sinh\left(\pi t\sqrt{1+s^2}\right)} \tag{3.66}$$

$$\varphi(t) = \varphi_{\widetilde{L}}(t)\,\varphi_{\widetilde{X}}(t) = e^{it\mu} \frac{\pi t}{\sinh(\pi t)} \frac{\pi t s}{\sinh(\pi t s)} \tag{3.67}$$

so that their difference is

$$\widetilde{\varphi}(t) - \varphi(t) = e^{it\mu}\left(\frac{\pi t\sqrt{1+s^2}}{\sinh\left(\pi t\sqrt{1+s^2}\right)} - \frac{\pi t}{\sinh(\pi t)}\frac{\pi t s}{\sinh(\pi t s)}\right). \tag{3.68}$$

Since $|e^{it\mu}| = 1$ for all real t and since the term in parentheses above is real,

$$\Delta(t, s) = \left|\widetilde{\varphi}(t) - \varphi(t)\right| = \frac{\pi t\sqrt{1+s^2}}{\sinh\left(\pi t\sqrt{1+s^2}\right)} - \frac{\pi t}{\sinh(\pi t)}\frac{\pi t s}{\sinh(\pi t s)}\,, \tag{3.69}$$

which is independent of μ. Note also that $x/\sinh(x)$ is symmetric about $x = 0$, so that $\Delta(t, s)$ will also be symmetric about $t = 0$ and we therefore only need to consider $t \geq 0$. To determine the maximum of Δ we plot it against t for different values of s, the results of which are shown in fig. 3.7. We notice that Δ is unimodal and falls quickly after its peak and it appears to attain its maximum when $s = 1$. To confirm this, rewrite Δ as

$$\Delta(t, s) = G(t\sqrt{1+s^2}) - G(t)G(ts)\,, \tag{3.70}$$

where $G(x) = \pi x/\sinh(\pi x)$. Then the maximiser $t^*(s)$ of Δ for a given s satisfies the first-order condition

$$\Delta_t(t^*(s), s) = \sqrt{1+s^2}G'(t^*\sqrt{1+s^2}) - sG(t^*)G'(t^*s) - G'(t^*)G(t^*s) = 0\,. \tag{3.71}$$

By the envelope theorem, the derivative with respect to s of maximum value of Δ is

$$\frac{d}{ds}\Delta(t^*(s), s) = \Delta_s(t^*(s), s) = \frac{t^*s}{\sqrt{1+s^2}}G'(t\sqrt{1+s^2}) - t^*G(t^*)G'(t^*s)\,, \tag{3.72}$$

so that the values t^* and s^* that maximise Δ for all t and s satisfy the two conditions

$$\sqrt{1+s^{*2}}G'(t^*\sqrt{1+s^{*2}}) = s^*G(t^*)G'(t^*s^*) - G'(t^*)G(t^*s^*) \tag{3.73}$$

$$\frac{t^*s^*}{\sqrt{1+s^{*2}}}G'(t\sqrt{1+s^{*2}}) = t^*G(t^*)G'(t^*s^*)\,. \tag{3.74}$$

Dividing eq. (3.73) by eq. (3.74) we have

$$\frac{1+s^{*2}}{s^*} = s^* + \frac{G(t^*)G(t^*s^*)}{G(t^*)G'(t^*s^*)} \tag{3.75}$$

$$\Longrightarrow\; G'(t^*s^*)G(t^*) = s^*G'(t^*)G(t^*s^*)\,, \tag{3.76}$$

which is satisfied by $s^* = 1$. The upper bound on Δ is then found by solving the first-order condition eq. (3.71) numerically (an analytical solution does not seem possible) when $s = 1$.■

Proof 3.B.3 (Proof of lemma 3.A.2) If $\alpha = 1$ and $\beta = 1$, then $f_X(t) = 1/(1+t)^2$ and we can find $P_X(x)$ by evaluating the integral in eq. (3.8) directly:

$$P_X(x) = \int_0^\infty \frac{x}{x+t}\frac{1}{(1+t)^2}\,dt \tag{3.77}$$

$$= \int_0^\infty \frac{x}{(x-1)^2}\left(\frac{x-1}{(1+t)^2} - \frac{1}{1+t} + \frac{1}{x+t}\right)dt \tag{3.78}$$

$$= \frac{x}{(x-1)^2}\left[-\frac{x-1}{1+t} + \frac{\ln(x+t)}{\ln(1+t)}\right]_0^\infty \tag{3.79}$$

$$= \frac{x(x-1-\ln x)}{(x-1)^2}\,. \tag{3.80}$$

To evaluate the integral we have used the partial fraction expansion

$$\frac{1}{(x+t)(1+t)^2} = \frac{A}{(1+t)^2} + \frac{B}{1+t} + \frac{C}{x+t} \tag{3.81}$$

$$= \frac{(Ax+Bx+c) + (A+B+Bx+2C)t + (B+C)t^2}{(x+t)(1+t)^2}\,, \tag{3.82}$$

$$= \frac{1}{(x-1)^2}\left(\frac{x-1}{(1+t)^2} - \frac{1}{1+t} + \frac{1}{x+t}\right), \tag{3.83}$$

obtained by setting $C = -B$ and $A = -B(x-1)$ to eliminate all the terms in t in the numerator of the right-hand side of eq. (3.82), which gives $B = -1/(x-1)^2$.

Differentiating $P_X(x)$ with respect to x gives $p_X(x)$. ■

Proof 3.B.4 (Proof of lemma 3.4) To show that f_V is unimodal we need to show that either f_V is decreasing for all $v \geq 0$ or that f_V increases and then decreases as v increases from 0. To do this, note that

$$q_V(p) = Q'_V(p) = (F_V^{-1})'(p) = \frac{1}{f_V(F_V^{-1}(p))} \iff f_V(v) = \frac{1}{q_V(F_V(v))} \tag{3.84}$$

so that

$$f_V'(v) = -\frac{1}{q(F_V(v))^2} f_V(v)\, q'(F_V(v)) = -\frac{q_V'(F_V(v))}{q_V(F_V(v))^3}\,. \tag{3.85}$$

Hence f_V is increasing iff the quantile density $q_V(v) = Q_V'(v)$ is increasing. Differentiating Q_V in eq. (3.49) twice with respect to p we can plot the sign of $q_V'(p)$ for all values of $p < 1$ (since Q_V is not differentiable for $p = 1$) and γ. Figure 3.10 shows that indeed q_V has one minimum over $p \in [0, 1)$, which means that f_V has one maximum, which is 0 for γ below a certain threshold.

Figure 3.10: Sign of $q_V'(p)$

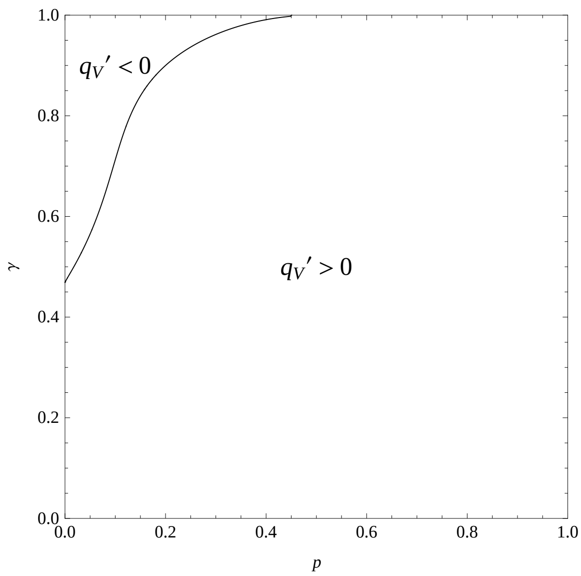

The median of V is easily computed by substituting $p = 1/2$ into the quantile function of V in eq. (3.49), yielding $4\alpha/\gamma$, as required.

To compute the mean of V, we use eq. (3.50) so that

$$\mathbb{E}(V) = \frac{\alpha}{\gamma}\left(\mathbb{E}\big(Y^{(1-\gamma)\tilde{\gamma}}\big) + 2\mathbb{E}\big(Y^{\tilde{\gamma}}\big) + \mathbb{E}\big(Y^{(1+\gamma)\tilde{\gamma}}\big)\right) \tag{3.86}$$

$$= \frac{\alpha}{\gamma}\left(\frac{\pi(1-\gamma)\tilde{\gamma}}{\sin(\pi(1-\gamma)\tilde{\gamma})} + 2\frac{\pi\tilde{\gamma}}{\sin(\pi\tilde{\gamma})} + \frac{\pi(1+\gamma)\tilde{\gamma}}{\sin(\pi(1+\gamma)\tilde{\gamma})}\right), \tag{3.87}$$

where we have used lemma 3.3 and the expression for the mean in definition 3.1. Now this expression is only valid if the powers inside the parentheses in eq. (3.86) are strictly less than 1. This occurs if the largest of the three, namely $(1+\gamma)\tilde{\gamma}$, is less than 1, i.e.,

$$(1+\gamma)\tilde{\gamma} = (1+\gamma)\sqrt{1-\gamma^2}/\gamma < 1 \iff (1+\gamma)^2(1-\gamma^2) < \gamma^2 \tag{3.88}$$

$$\iff \gamma^4 + 2\gamma^3 + \gamma^2 - 2\gamma - 1 > 0\,. \tag{3.89}$$

Figure 3.11 shows that for $\gamma \in (0, 1)$, this holds if γ is larger than a threshold γ^*, which is about 0.883.

Figure 3.11: Plot of $\gamma^4 + 2\gamma^3 + \gamma^2 - 2\gamma - 1$

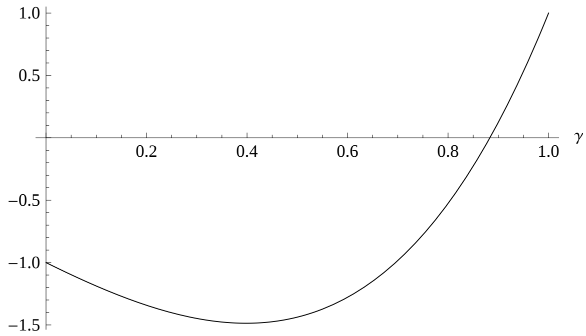

Proof 3.B.5 (Proof of lemma 3.6) Theorem 3 in Wilfling (1996) states that if $g, h : \mathbb{R}^+ \to \mathbb{R}^+$ are positive, continuous and increasing transformations and $g(x)/h(x)$ is decreasing on $\mathbb{R}^+$, then $g(X) \leq_L h(X)$ if $g(X)$ and $h(X)$ have finite expectation.

Since the Lorenz order is scale-invariant, without loss of generality, we can normalise the scale parameter to 1. Then, if $V_1 \sim \mathrm{VL}(1, \gamma_1)$ and $V_2 \sim \mathrm{LL}(1, \gamma_2)$, by eq. (3.50), $V_i \overset{d}{=} g_i(Y)$ where

$$g_i(x) = \frac{1}{\gamma_i}\left(x^{(1-\gamma_i)\widetilde{\gamma}_i} + 2x^{\widetilde{\gamma}_i} + x^{(1+\gamma_i)\widetilde{\gamma}_i}\right) \tag{3.90}$$

$$= \frac{1}{\gamma_i} x^{(1-\gamma_i)\widetilde{\gamma}_i}\left(1 + 2x^{\gamma_i\widetilde{\gamma}_i} + x^{2\gamma_i\widetilde{\gamma}_i}\right) \tag{3.91}$$

$$= \frac{1}{\gamma_i} x^{(1-\gamma_i)\widetilde{\gamma}_i}\left(1 + x^{\gamma_i\widetilde{\gamma}_i}\right)^2. \tag{3.92}$$

Then

$$\frac{g_1(x)}{g_2(x)} = \frac{\gamma_2}{\gamma_1} x^{(1-\gamma_1)\widetilde{\gamma}_1-(1-\gamma_2)\widetilde{\gamma}_2}\left(\frac{1 + x^{\gamma_1\widetilde{\gamma}_1}}{1 + x^{\gamma_2\widetilde{\gamma}_2}}\right)^2. \tag{3.93}$$

Since $\gamma_1 > \gamma_2 \iff \widetilde{\gamma}_1 < \widetilde{\gamma}_2$, $\gamma_1 > \gamma_2 \iff (1-\gamma_1)\widetilde{\gamma}_1 < (1-\gamma_2)\widetilde{\gamma}_2$ and $\gamma_1 > \gamma_2 \iff \gamma_1\widetilde{\gamma}_1 < \gamma_2\widetilde{\gamma}_2$, so that $g_1(x)/g_2(x)$ is monotonically decreasing iff $\gamma_1 > \gamma_2$, so that applying Theorem 3 in Wilfling (1996) gives the result. ■

Proof 3.B.6 (Proof of proposition 3.6) Using eq. (3.29) we know that the sign of $\partial_\gamma x^*(v; \alpha, \gamma)$ is the same as that of

$$\partial_\gamma p_X(x; \alpha, \gamma) = \frac{\partial}{\partial\gamma}\frac{1}{x}\frac{\gamma\,(x/\alpha)^\gamma}{(1 + (x/\alpha)^\gamma)^2} \tag{3.94}$$

$$= \frac{(x/\alpha)^\gamma}{x(1 + (x/\alpha)^\gamma)^3}\left((x/\alpha)^\gamma + 1 - \left((x/\alpha)^\gamma - 1\right)\gamma\log(x/\alpha)\right) \tag{3.95}$$

so that $\partial_\gamma x^*(Q_X(p); \alpha, \gamma)$, the sign of $\partial_\gamma x^*$ at the p-th quantile of X, is equal to the sign of

$$S(p; \gamma) := \left(\frac{p}{1-p}\right)^{\sqrt{1-\gamma^2}} + 1 - \left(\left(\frac{p}{1-p}\right)^{\sqrt{1-\gamma^2}} - 1\right)\sqrt{1-\gamma^2}\log\left(\frac{p}{1-p}\right), \tag{3.96}$$

which is obtained by substituting the expression for Q_X in eq. (3.32). Note that $S(p;\gamma)$ is symmetric in p about $p = 1/2$. To see this, note that

$$S(1-p;\gamma) = \left(\frac{1-p}{p}\right)^{\sqrt{1-\gamma^2}} + 1 - \left(\left(\frac{1-p}{p}\right)^{\sqrt{1-\gamma^2}} - 1\right)\sqrt{1-\gamma^2}\log\left(\frac{1-p}{p}\right) \tag{3.97}$$

$$= 1 + \left(\frac{p}{1-p}\right)^{\sqrt{1-\gamma^2}} + \left(1 - \left(\frac{p}{1-p}\right)^{\sqrt{1-\gamma^2}}\right)\sqrt{1-\gamma^2}\log\left(\frac{p}{1-p}\right) \tag{3.98}$$

$$= S(p;\gamma)\,. \tag{3.99}$$

Therefore, the sign of $\partial_\gamma x^*$ is the same for the top and bottom quantiles of effort and therefore values. Also, note that $S(1/2;\gamma) = 2 > 0$ for all γ, so $\partial_\gamma x^*(Q_X(p)) > 0$ for all $p \in (1/2 - \delta(\gamma), 1/2 + \delta(\gamma))$, where $p = 1/2 - \delta(\gamma)$ and $p = 1/2 + \delta(\gamma)$ are the two solutions to $S(p;\gamma) = 0$. To see how δ changes with γ, note that S is a function of p and γ only through $(p/(1-p))^{\sqrt{1-\gamma}}$, so that to keep $S(p,\gamma) = 0$ when γ increases, $p/(1-p)$ and therefore p must decrease. ∎

Proof 3.B.7 (Proof of proposition 3.7) Firstly, in order for the expectation of V to be well-defined, we need $\gamma > \gamma^*$. In that case, letting $\mu = \mathbb{E}V$, the scale parameter α must satisfy

$$\alpha = \mu/H(\gamma) = \mu\gamma\left(\frac{\pi(1-\gamma)\widetilde{\gamma}}{\sin(\pi(1-\gamma)\widetilde{\gamma})} + 2\frac{\pi\widetilde{\gamma}}{\sin(\pi\widetilde{\gamma})} + \frac{\pi(1+\gamma)\widetilde{\gamma}}{\sin(\pi(1+\gamma)\widetilde{\gamma})}\right)^{-1}, \tag{3.100}$$

so that

$$\frac{\mathbb{E}X}{\mathbb{E}V} = \alpha/\mu\frac{\pi\widetilde{\gamma}}{\sin(\pi\widetilde{\gamma})} = \gamma\frac{\pi\widetilde{\gamma}}{\sin(\pi\widetilde{\gamma})}\left(\frac{\pi(1-\gamma)\widetilde{\gamma}}{\sin(\pi(1-\gamma)\widetilde{\gamma})} + 2\frac{\pi\widetilde{\gamma}}{\sin(\pi\widetilde{\gamma})} + \frac{\pi(1+\gamma)\widetilde{\gamma}}{\sin(\pi(1+\gamma)\widetilde{\gamma})}\right)^{-1}. \tag{3.101}$$

Since eq. (3.101) is a function of one variable over a bounded interval, rather than differentiating with respect to γ, we can simply plot it to show that it's monotonically increasing, as in fig. 3.5. ∎

Proof 3.B.8 (Proof of proposition 3.8) If we multiply eq. (3.15) by x^*, we have

$$\frac{x^*(v)}{v} = x^*(v)\, p_X(x^*(v)) = \frac{\gamma(x^*(v)/\alpha)^\gamma}{(1+(x^*(v)/\alpha)^\gamma)^2}\,, \tag{3.102}$$

which indeed hump-shaped since its derivative with respect to v is

$$x^{*\prime}(v)(x^*(v)/\alpha)^{\gamma-1}\frac{\gamma(1-(x^*(v)/\alpha)^\gamma)}{(1+(x^*(v)/\alpha)^\gamma)^3}\,, \tag{3.103}$$

which is positive if $x^*(v) < \alpha$ and negative if $x^*(v) > \alpha$, so that it has a maximum of $\gamma/4$ at median effort and hence median value. We can also write it as

$$w_X(p) = w_X(Q_X(p)) = \frac{\gamma \left(\frac{p}{1-p}\right)^{\sqrt{1-\gamma^2}}}{\left(1 + \left(\frac{p}{1-p}\right)^{\sqrt{1-\gamma^2}}\right)^2}. \tag{3.104}$$

Now

$$\frac{d}{d\gamma} \frac{\left(\frac{p}{1-p}\right)^{\sqrt{1-\gamma^2}}}{\left(1 + \left(\frac{p}{1-p}\right)^{\sqrt{1-\gamma^2}}\right)^2} = \frac{\gamma \left(\frac{p}{1-p}\right)^{\sqrt{1-\gamma^2}}}{\left(1 + \left(\frac{p}{1-p}\right)^{\sqrt{1-\gamma^2}}\right)^3} \frac{\gamma \left(\left(\frac{p}{1-p}\right)^{\sqrt{1-\gamma^2}} - 1\right) \log\left(\frac{p}{1-p}\right)}{\sqrt{1-\gamma^2}}, \tag{3.105}$$

which is always non-negative, since $(p/(1-p))^{\sqrt{1-\gamma^2}} - 1$ and $\log(p/(1-p))$ always have the same sign. Hence $w_X(p)$ is non-decreasing in γ. ■

Bibliography

Acemoglu, D. (2003). “Why not a Political Coase Theorem? Social Conflict, Commitment, and Politics”. *Journal of Comparative Economics*, vol. 31.

Acemoglu, D., G. Egorov, and K. Sonin (2009). “Do Juntas Lead to Personal Rule?” *American Economic Review: Papers and Proceedings*, vol. 99, no. 2, pp. 298–303.

Acemoglu, D., S. Johnson, and J. A. Robinson (2001). “The Colonial Origins of Comparative Development: An Empirical Investigation”. *American Economic Review*, vol. 91, no. 5.

Acemoglu, D. and J. Robinson (2006). *Economic Origins of Dictatorship and Democracy*. MIT press.

Allingham, M. G. and A. Sandmo (1972). “Income tax evasion: A theoretical analysis”. *Journal of Public Economics*, vol. 1, pp. 323–338.

Alm, J. and M. McKee (2004). “Tax compliance as a coordination game”. *Journal of Economic Behavior and Organization*, vol. 54, pp. 297–312.

Athey, S. (2002). “Monotone comparative statics under uncertainty”. *Quarterly Journal of Economics*, vol. 117, pp. 187–223.

Atkinson, A. B. (1970). “On the Measurement of Inequality”. *Journal of Economic Theory*, vol. 2, pp. 244–263.

Besley, T. and M. Ghatak (2010). “Property Rights and Economic Development”. In: ed. by D. Rodrik and M. Rosenzweig. North-Holland. Chap. 68.

Besley, T. and J. A. Robinson (2009). “Quis Custodiet Ipsos Custodes? Civilian Control over the Military”. unpublished.

Billingsley, P. (1995). *Probability and Measure*. 3rd ed. Wiley.

Dal Bó, E. (2007). “Bribing Voters”. *American Journal of Political Science*, vol. 51, no. 4, pp. 789–803.

Ewerhart, C. (2010). *Rent-seeking contests with independent private values*. Working Paper 490. University of Zurich.

Fearon, J. D. (1995). “Rationalist explanations for war”. *International Organization*, vol. 49, no. 3, pp. 379–414.

Feldman, M. and C. Gillies (1985). “An expository note on individual risk without aggregate uncertainty”. *Journal of Economic Theory*, vol. 39, pp. 26–32.

Fey, M. (2008). “Rent-seeking contest with incomplete information”. *Public Choice*, vol. 135, pp. 225–236.

FTC (2011). “IRS analyses and summaries as investigatory records.” In: *Federal Tax Coordinator*. T-10308.

GAO (2001). *IRS Audit Rates: Rate for Individual Taxpayers Has Declined But Effect on Compliance Is Unknown*. Report to the Chairman, Subcommittee on Oversight, Committee on Ways and Means, House of Representatives GAO-01-484. United States General Accounting Office.

Garfinkel, M. and S. Skaperdas (2007a). "Economics of conflict". In: *Handbook of Defense Economics*. Vol. 2. North-Holland/Elsevier, pp. 649–709.

Garfinkel, M. and S. Skaperdas (2007b). "Economics of Conflict: An Overview". In: ed. by K. Hartley and T. Sandler. Handbook of Defense Economics.

Grossman, H. I. and M. Kim (1995). "Swords or Plowshares? A Theory of the Security of Claims to Property". *Journal of Political Economy*, vol. 103, no. 6, pp. 1275–1288.

Grossman, H. I. and S. J. Noh (1990). "A Theory of Kleptocracy with Probabilitic Survival and Reputation". *Economics and Politics*, vol. 2, no. 2.

Guiso, L., P. Sapienza, and L. Zingales (2006). "Does Culture Affect Economic Outcomes?" *Journal of Economic Perspectives*.

Hansen, L. G., S. Krarup, and C. S. Russell (2006). "Enforcement and information strategies". *Journal of Regulatory Economics*, vol. 30, pp. 45–61.

Hillman, A. and J. Riley (1989). "Politically contestable rents and transfers". *Economics and Politics*, vol. 1, pp. 17–39.

Hirshleifer, J. (1995). "Anarchy and its Breakdown". *Journal of Political Economy*, vol. 103, no. 1, pp. 26–52.

Hornblower, S. and A. Spawforth, eds. (2003). *The Oxford Classical Dictionary*. 3rd ed. Oxford University Press.

Hurley, T. and J. Shogren (1998). "Asymmetric information contests". *European Journal of Political Economy*, vol. 14, pp. 645–665.

Jones, L. (2000–2001). "Strange Words – Audit and Examination". *Journal of Tax Practice and Procedure*, vol. 2, pp. 5–6.

Judd, K. L. (1985). "The law of large numbers with a continuum of iid random variables". *Journal of Economic Theory*, vol. 39, pp. 19–25.

Kleiber, C. and S. Kotz (2003). *Statistical Size Distributions in Economics and Actuarial Sciences*. Wiley.

Knack, S. and P. Keefer (1995). "Institutions and Economics Performance: Cross-Country Tests Using Alternative Measures". *Economics and Politics*, vol. 7, no. 3.

Macho-Stadler, I. and J. D. Perez-Castrillo (2002). "Auditing with signals". *Economica*, vol. 69, pp. 1–20.

Malueg, D. and A. Yates (2004). "Rent seeking with private values". *Public Choice*, vol. 119, pp. 161–178.

McGuire, M. C. and M Olson (1996). "The economics of autocracy and majority rule: The invisible hand and the use of force". *Journal of Economic Literature*, vol. 34, no. 1, pp. 72–96.

McNeil, A. J., R. Frey, and P. Embrechts (2005). *Quantitative Risk Management*. Princeton University Press.

Milgrom, P. R. (1981). “Good News and Bad News: Representation Theorems and Applications”. *Bell Journal of Economics*, vol. 12, no. 2, pp. 380–391.

Morris, S. and H. Shin (2003). “Global Games: Theory and Applications”. In: *Advances in Economics and Econometrics: Theory and Applications*. Ed. by M. Dewatripont, L. P. Hansen, and S. J. Turnovsky. Vol. 1. Proceedings of the Eighth World Congress of the Econometric Society. Cambridge University Press, pp. 56–114.

Nelsen, R. B. (2006). *An Introduction to Copulas*. 2nd ed. Springer.

North, D. C., J. J. Wallis, and B. R. Weingast (2005). “The Natural State: The Political-Economy of Non-Development”. *Unpublished*.

Ok, E. A. (2011). “Expectation via the Stieltjes Integral”. In: *Probability Theory with Economic Applications*. *unpublished*. Chap. D.

Olson, M. (1993). “Dictatorship, Democracy, and Development”. *American Political Science Review*, vol. 87, no. 3, pp. 567–576.

Ormiston, M. B. and E. E. Schlee (1993). “Comparative Statics under Uncertainty for a Class of Economic Agents”. *Journal of Economic Theory*, vol. 61, pp. 412–422.

Piccione, M. and A. Rubinstein (2007). “Equilibrium in the Jungle”. *Economic Journal*, vol. 117, no. 522, pp. 883–896.

Prada-Sarmiento, J. D. (2010). *Uncertainty in Conflicts*. Documentos CEDE. Universidad de los Andes, CEDE.

Reinganum, J. F. and L. L. Wilde (1985). “Income tax compliance in a principal-agent framework”. *Journal of Public Economics*, vol. 26, no. 1, pp. 1–18.

— (1986). “Equilibrium Verification and Reporting Policies in a Model of Tax Compliance”. *International Economic Review*, vol. 27, no. 3, pp. 739–760.

— (1988). “A Note on Enforcement Uncertainty and Taxpayer Compliance”. *Quarterly Journal of Economics*, vol. 103, no. 4, pp. 793–798.

Roberts v. IRS (1984). *543 Federal Supplement 1241*.

Ryvkin, D. (2010). “Contests with private costs: Beyond two players”. *European Journal of Political Economy*, vol. 26, pp. 558–567.

Sanchez, I. and J. Sobel (1993). “Hierarchical design and enforcement of income tax policies”. *Journal of Public Economics*, vol. 50, no. 3, pp. 345–369.

Sanchez-Villalba, M. (2006). “Anti-evasion auditing policy in the presence of common income shocks”. Working paper.

Skaperdas, S. (1992). “Cooperation, Conflict, and Power in the Absence of Property Rights”. *American Economic Review*, vol. 82, no. 4, pp. 720–739.

Sun, Y. (1998). “A theory of hyperfinite processes: the complete removal of individual uncertainty via exact LLN”. *Journal of Mathematical Economics*, vol. 29, pp. 419–503.

— (2006). “The exact law of large numbers via Fubini extension and characterization of insurable risks”. *Journal of Economic Theory*, vol. 126, pp. 31–69.

Tan, F. and A. Yim (2011). *Can transparency hurt? An experiment on whether disclosure of audit policy details reduces tax compliance*. MPRA Paper 31160.

Taylor, L. R. (2003). *Roman Voting Assemblies: From Hannibalic War to the Dictatorship of Caesar*. The University of Michigan Press.

Tullock, G. (1980). "Efficient rent seeking". In: *Toward a Theory of Rent-Seeking Society*. Ed. by J. Buchanan, R. Tollison, and G. Tullock. Texas A&M University Press, pp. 97–112.

Wärneryd, K. (2010). "Informational Aspects of Conflict". In: *Oxford Handbook of the Economics of Peace and Conflict*. Ed. by M. Garfinkel and S. Skaperdas. Oxford University Press. Chap. 2.

Wasser, C. (2011). "Incomplete Information in Rent-seeking Contests". *Economic Theory*. (forthcoming).

Wilfling, B. (1996). "Lorenz ordering of generalized beta-II income distributions". *Journal of Econometrics*, vol. 71, pp. 381–388.