Introduction

This introduction aims at presenting an overview of the content of the book. For each chapter it presents the problem, the model and the main results. Whenever possible, an analysis or even a sketch of the proof is given within an appropriately simplified version of the model. The information can also be read by section, each section, or overview, acting as an extended introduction to the corresponding chapter. Finally, it provides a summary to which the reader may want to refer after reading the whole book. Naturally, these are not mutually exclusive uses of the Introduction!

As stressed in the foreword, and as emphasized in the title, the monograph is a contribution to pure theory. As such, it aims at exploring the logic of an abstract model, the stylized features of which describe a (so-called) polar world. Within such a model – again 'théorie pure' à la Walras – derivations and also comments concern the polar world under consideration and not the specific issues, debates, and controversies that the analysis of the polar world aims at clarifying. Although the separation of theory and policy analysis is a standard modern procedure, it has the inconvenience of making the theorist's logic and motivation more obscure to practitioners than it should be. As an attempt to overcome this difficulty, at the beginning of each overview a subsection, entitled 'Motivation of the chapter: issues, studies and debates', sketches a description of background materials, underlying issues, related applied studies and possible controversies that put each chapter in a better perspective. Throughout the book bibliographical references are kept to a minimum. In fact, they are gathered as much as possible in the bibliographical note appearing at the end of each chapter.

1 An overview of chapter 1: the institutional economics of taxation

1.1 Motivation of the chapter: related issues, studies and debates

Chapter 1 mainly attempts to answer a question that is primarily theoretical. To what extent and under what conditions are taxation schemes of the kind commonly observed – consumption taxes such as VAT, income tax, etc. – appropriate taxing devices?

This question has a long history: an influential segment of the profession has argued that personalized transfers perform better than distortionary taxes. This issue still creeps into current debates on policy (it will for example be referred, in the introduction to chapter 1, to the French debate on housing subsidies in the 1970s). It can however be put into better perspective by the modern theory of incentives. This theory puts at the forefront the informational problems of public organizations and suggests for example the reconsideration of the controversies surrounding the relative merits of lump-sum taxation and distortionary taxes in the light of these informational problems. This is what chapter 1 is aimed at.

To some extent, the basic message gives support to a reasonable intuition: that the taxation base may include all the (easily) observable variables or actions of individuals but need not include too complex considerations – such as making somebody's taxes depend on somebody else's taxes or depend on a complex set of announcements. In this sense, the standard tax system of the theory of second best, that mixes linear taxes and non-linear taxes, can be justified by incentive theory arguments, those that relate to the so-called taxation principle.

This message, satisfactory from the intuitive or practioner's viewpoint, has however to be reconciled with the fact that incentives theory has designed powerful mechanisms that seem to perform better, at least in some contexts, than standard taxation devices. It is shown, at the end of chapter 1, that one key issue is correlation. More complex incentives devices will not perform better when the agents' 'hidden' characteristics – those that would be fully relevant for taxation purposes – are uncorrelated. Indeed, the most spectacular tools of incentive theory – those associated with the so-called Nash, perfect Nash or even Bayesian Nash implementation – exploit to a considerable extent correlation of the information held by individual agents.

Hence chapter 1 attempts to base a theory of tax institutions on informational and incentive considerations. It does not claim to be exhaustive: the actual tax base should depend on the objectives of the 'government', i.e., on optimization (see chapter 3), but also on the 'administrative' costs of including such-or-such variable in the tax base (in the present analysis these costs are zero or plus infinity). However chapter 1, besides providing a background to the recent contributions in public finance that stress self-selection problems[1] (particularly for studying anti-poverty mechanisms) presents a starting point for a more comprehensive

[1] See for example Blackorby–Donaldson (1988), Besley–Coates (1992).

look at the tax institutions (such as the one sketched in the conclusion of this book).

1.2 The methodology from a simple model

The aim of chapter 1 is to provide one (not necessarily the only one) coherent justification of the model that will be considered later. This model is a taxation model in which commodity and factor taxes introduce wedges between the prices faced by the production sector and the prices faced by the consumption sector. An obvious objection to such taxation schemes is that they distort the choice of economic units and create inefficiencies. From the second welfare theorem, these inefficiencies could be removed if distorting taxes were replaced by lump-sum transfers. To put it in another way, the model can be justified only (i) if lump-sum transfers are unavailable, (ii) if commodity taxes are the right substitutes for lump-sum taxes. Chapter 1 indeed argues that there are circumstances where for some basic reasons conditions (i) and (to some extent) (ii) are fulfilled.

The following discussion of a simple model provides an introduction to the core of the argument of chapter 1.

Let us consider here a two-good economy. Commodity 1 is labour, commodity 2 is the 'consumption good'. The economy has a large number of final agents (consumers) which we represent as a continuum of agents. Each agent has preferences depending upon a one-dimensional characteristic described by a parameter $\theta (\theta \in [0,1])$. Then, household θ's preferences over the bundle consumption–labour (c,l) are represented by a concave utility function $u(c,l,\theta)$. At this stage of the analysis, the exact nature of θ and l is not essential. For example, θ might be a personal productivity parameter,[2] which multiplied by the labour time effectively supplied (effective labour) would determine the quantity l of 'efficient labour'. The production sector is in charge of constant returns to scale techniques which, in the setting just sketched, transform one unit of labour into one unit of consumption. Let $(c(\theta), l(\theta))$ be the consumption – labour bundle of θ. If there is a continuum of characteristics distributed as described by some probability measure μ, then the scarcity constraint can be written

$$\int c(\theta) d\mu \leq \int l(\theta) d\mu \qquad (1)$$

Again, the fact that the individual utility function depends upon θ may describe differences in tastes but may also reflect – as in the just sketched productivity setting – the fact that agents of different productivities find it

[2] As in Mirrlees (1971).

unequally difficult to supply a given amount of 'efficient labour' (since this amount involves different numbers of hours of effective labour time).

Consider a utilitarian planner, whose social optimization problem would consist in maximizing $\int u(c(\theta), l(\theta), \theta)d\mu$ under the scarcity constraint(s).

The solution consists of a vector function that associates to every θ a bundle consumption – (efficient) labour. It is a **first-best optimum** in the sense that it is only constrained by scarcity of resources and technological limitations.

In fact, it is well known – since this follows from the second welfare theorem – that the optimum can be decentralized, i.e., can obtain as the outcome of a market organization. This market solution necessarily has the following features. First, the market price of the consumption good being set equal to one, the wage (of 'efficient labour') is also one. Equivalently in the productivity interpretation of the model the wage of 'efficient labour' is one and the wage of a unit of effective labour equals θ, i.e., its productivity. Second, the income of every agent does not only consist of his labour income but also incorporates some income transfer; here income transfer R is naturally indexed on θ and then necessarily satisfies $\int R(\theta)d\mu = 0$ (the net total transfer is zero).

As is well known, the basic message behind this 'simple' story has a broad validity: it holds true in a world with many commodities and with public goods: the attainment of the first-best optimum does not require and is in general incompatible with distorting taxes.

The simple story however makes it clear that the decentralization of the first-best solution requires the modification of the primary income distribution determined by market forces. Transfers, positive or negative that depend upon θ, have to be implemented. Their implementation requires that either the planner knows θ or, if it is not the case, that the agents – having information on their own θ – are willing to transmit it truthfully. In other words, the implementation of the first-best optimum requires either that the information on which lump-sum transfers are based is public information or, in the case where it is not public information, that it can be costlessly acquired, for example through an appropriate **incentive compatible device**.

We have not yet introduced the tools that allow the analysis of incentive compatibility. However we may already note that the revelation of the information on θ – assuming it is private – is unlikely. For example, if the above productivity interpretation is specified in such a way that the initial preferences between consumption and effective labour of all agents are assumed to be similar – although utility functions that trade off between consumption and efficient labour are different – then it can be shown that not only is taxation redistributive (R decreases when θ increases) but also that the welfare finally achieved by the agents is negatively correlated with their productivity.[3] In such a case it is easy to understand that highly productive agents have little incentive to confess their productivity!

1.3 The core of the argument

First-best organization of the economy would rule out commodity taxes. Informational problems make such a first-best infeasible and will force us to rely on taxes. To see that, let us come back to the above simple world. But let us explicitly assume that information on θ instead of being public – as the reference to the second welfare theorem implicitly supposes – is private: each agent knows its own characteristic. Let us suppose however that the centre – the planner – knows the probability distribution of characteristics.

In this setting, economic organization must take into account not only the technological and scarcity constraints but also the **informational constraints**. Its study should be viewed as a subject of the so-called theory of incentives.[4] In the light of this theory, let us consider incentive mechanisms that can be used by a planner. Let us restrict attention to a special kind of such incentive mechanisms, those that are direct, anonymous and truthful. A mechanism is direct when the agent's announcements only consists of his own characteristics instead of more abstract (and possibly complex) messages; a mechanism is anonymous when it depends upon the announcement of the agent and the distribution of the announcements of the others; a mechanism is truthful when it is designed in such a way that it is in each agent's interest to truthfully reveal his characteristics.

The reader may inquire about the significance of the restriction that the consideration of the above class of mechanisms induces. A complete answer cannot be given without an exposition of the theory of incentives that is out of the scope of this introduction (or even of this book). To shorten the story, the restriction to truthful direct mechanisms is innocuous in our setting: this is the celebrated revelation principle (that holds given the informational assumptions – which themselves affect our choice of solution concept) that we make more precise in chapter 1. Also, the study of the limitations induced by anonymity is discussed in chapter 1 (particularly section 1.6).

Formally, a **direct anonymous truthful mechanism** consists of a mapping F

[3] This (insufficiently known) result is due to Mirrlees (1986); see also Guesnerie (1980).

[4] Following Gibbard (1973), Hurwicz (1973), economic organization should be analysed from the study of 'game forms'. One important ingredient for the study of game forms is the choice of a solution concept. The present conventional wisdom in incentive theory is that the choice of the solution concept (dominant strategy, Bayesian–Nash, or Nash) should reflect the basic informational assumptions that are made (agents ignore the others' characteristics, know it statistically, know it exactly). For obvious reasons, this introduction remains rather loose on these questions.

which associates a vector of consumption and (efficient) labour as a function of an announcement of characteristics and of a distribution of announcements

$$F:(\theta,\nu)\in\Theta\times\mathscr{P}(\Theta)\to F(\theta,\nu)=\{F_1(\theta,\nu),F_2(\theta,\nu)\}\in\mathscr{R}^2$$

(where $\mathscr{P}(\Theta)$ designates the set of probability distributions over Θ) and which satisfies the following condition

$$\theta\in\arg\max_{\hat{\theta}} u(F(\hat{\theta},\nu),\theta),\forall\theta,\text{ for every }\nu \tag{2}$$

Condition (2) formally expresses that truthful revelation will obtain whatever the agent's characteristic, but also, here, whatever the announcement of others: in other words θ is a dominant strategy. Then we refer to the concept of dominant strategy implementation (in chapter 1 a less demanding Bayesian–Nash concept of implementation will be adopted).

Furthermore the mechanism F is **admissible** whenever it satisfies

$$\int_\Theta F_1(\theta,\nu)d\nu=\int_\Theta F_2(\theta,\nu)d\nu \tag{3}$$

i.e., when it induces the equality of total consumption with total supply of efficient labour. This equality is supposed to hold whatever the distribution of characteristics ν; in fact, the true distribution μ will necessarily be discovered so that it is enough, in a sense, to have (3) met with $\nu=\mu$. For that reason, we will often forget (here) about the dependence of ν on the mechanism.[5]

In this setting, the planner's problem is drastically different from the first-best problem: it is a second-best problem that consists of maximizing social welfare over the set of mechanisms that are truthful and admissible, i.e., that satisfy (2) and (3).

There is however a less sophisticated way to allocate resources when the information on characteristics is private: it is to rely upon **taxation systems**.

Here, l is the amount of efficient labour: if labour is valued in a competitive setting and if the price of the consumption good is one, l is identical to labour income. Assuming that l is observable – as it has been implicitly assumed in the definition of the above mechanism – the planner can consider implementing an income tax schedule.

Formally an **income tax schedule** is a mapping ψ which, for every positive labour income (pre-tax income), associates some (positive) consumption level (post-tax income).

Given the income tax schedule ψ, a household with productivity θ

[5] However in the terminology of incentive theory, 'admissibility' will only hold at equilibrium and not outside equilibrium.

chooses his bundle consumption–efficient labour $c_\psi(\theta)$, $l_\psi(\theta)$ as a solution of the following programme $P(\theta,\psi)$

$$\max u(c,l,\theta) \qquad (4)$$
$$c \leq \psi(l)$$

The tax schedule ψ is admissible if

$$\int_\Theta c_\psi(\theta)d\mu = \int_\Theta l_\psi(\theta)d\mu \qquad (5)$$

i.e., if total consumption equals total labour supply.

The planner's problem when he considers the implementation of income tax schedules is still a **second-best problem**: to maximize social welfare over the set of admissible income tax schedules.

It is natural now to look at a **sophisticated planner** – or the incentive mechanism planner – and an **unsophisticated** one – or the tax schedule planner – and to compare their performances.

The comparison is strikingly simple: the set of allocations that can be achieved by, on the one hand, the sophisticated and, on the other hand, the unsophisticated planner are **the same** (and consequently both can achieve the same level of social welfare).

The proof of this **equivalence result** is simple and can be sketched as follows:

(i) Consider some tax schedule ψ and let $c_\psi(\theta)$, $l_\psi(\theta)$ be one corresponding outcome (obtained as the solution of programme (4) $P(\theta,\psi)$). Then the mapping $\theta \rightarrow c_\psi(\theta)$, $l_\psi(\theta)$ defines a direct anonymous truthful mechanism that is admissible when the true distribution is μ.[6] This statement is easy to check. Assume that it is wrong, i.e., that there exists θ and $\hat{\theta}$ such that $c_\psi(\hat{\theta})$, $l_\psi(\hat{\theta})$ is better for household θ than $c_\psi(\theta)$, $l_\psi(\theta)$; but $c_\psi(\hat{\theta})$, $l_\psi(\hat{\theta})$ is a point on the income tax schedule – the one chosen by $M.\hat{\theta}$ – that would have been preferred by household θ to $c_\psi(\theta)$, $l_\psi(\theta)$ in the solution of programme (4), a contradiction.

(ii) Consider now some truthful incentive mechanism $c(\theta,\mu)$, $l(\theta,\mu)$ (restricted to the given distribution μ). We can show that there is an income tax schedule ψ that leads to the same outcome.

For that it is enough to consider $Z = U_\theta\{c(\theta,\cdot), l(\theta,\cdot)\}$ a subset of $\mathscr{R}^2_+$ and then the north-east frontier of the set $Z - \mathscr{R}^2_+$. This frontier is the graph of a

[6] In the spirit of the previous remarks, we focus attention on the actual distribution μ. Naturally, the analysis of chapter 1 is much more explicit on the significance of the choice of what incentive theory terms the 'domain of environments'.

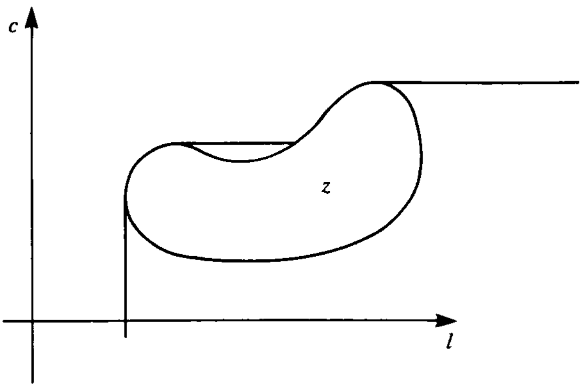

Figure 1

function ψ (Figure 1). Arguing as above by contradiction, one concludes that a solution of household θ's maximization problem when it faces the income tax schedule ψ is indeed $c(\theta,\cdot)l(\theta,\cdot)$.[7]

The conclusion is clearcut: 'sophisticated' incentive mechanisms can do no better in our problem than 'unsophisticated' income tax schedules. (Note however that although unsophisticated our tax schedules are not linear and can only be approximated by linear tax schedules.)

1.4 Developments of the argument

The preceding equivalence result is at the core of the argument of the first part of chapter 1; this argument is however more intricate for at least two reasons. On the one hand, one considers an n-commodity world instead of a two-commodity model; on the other hand, one distinguishes two kinds of private commodities: the first ones cannot be exchanged by households once they have been allocated by the planner, but the second ones are outside the planner's control. Consequently the unsophisticated planner is restricted to using linear taxation for the 'tradable' goods, keeping non-linear taxation for the 'non-tradable' ones.

In the just sketched setting chapter 1 establishes an equivalence result that generalizes the equivalence result briefly described here. Again the sophisticated planner, using direct anonymous truthful admissible mechanisms subject to the tradability constraint, and the unsophisticated planner using tax systems that are partly linear and partly non-linear, can implement the same final allocations.

[7] Such a property has a general counterpart, called by Rochet (1986) the taxation principle, see Hammond (1979) and Guesnerie (1981a).

The equivalence result is however conditional on two key ingredients of the analysis that have to be discussed further.

First, the equilibrium concept adopted for the analysis of the game form is the concept of Bayesian–Nash equilibrium.

This condition is clearly restrictive: adopting the concept of a perfect-Nash equilibrium, recent work[8] that applies to the present setting has argued that 'almost every performance function' could be implemented. There is however no contradiction between this conclusion and that of chapter 1: Nash implementation makes complete sense only when agents know their own characteristics as well as the characteristics of other agents. The concept of a Bayesian–Nash equilibrium adopted here reflects a more restrictive informational hypothesis, i.e., that each agent knows his own characteristics and only the distribution of characteristics of others.

Second, the sophisticated planner is supposed to use direct anonymous truthful mechanisms rather than mechanisms that will rely upon a more complex set of messages and/or will not be anonymous.

As mentioned, a general principle of incentive theory, the revelation principle, holds (for Bayesian–Nash implementation) and indicates that there is no loss of generality in taking direct truthful mechanisms. There is however a problem with anonymity: it is discussed at length in the last section of chapter 1. The main conclusion is that the kind of anonymity assumed here involves no loss of generality for the analysis of incentives only when there is **no correlation between the individual characteristics of agents**: this obtains if for example each agent's characteristics are the outcomes of independent drawings from the same basket. In that case the profile of characteristics is the realization of an infinite number of independent and identically distributed random variables and a variant of the equivalent result emphasized holds true here.

The above informal analysis gives the appropriate flavour of the content of chapter 1. However the intuition that is conveyed here cannot be made rigorous without some precautions: this involves distinguishing the names of agents from their characteristics, making assumptions of how goods are traded outside the planner's control, making explicit the implicit assumptions on observability etc. It follows that the reading of chapter 1 requires more caution than its informal summary suggests.

In conclusion chapter 1 establishes general conditions under which the kind of organization described in the model of this book – which we associated with unsophisticated planning – is the best that can be achieved. However, the model considered in this book only considers linear taxation – or 'affine' taxation when a uniform lump sum is considered – when the equivalence result supports the use of non-linear taxation for a subset of

[8] See for example Moore (1992).

commodities. In other words, the model of this book assumes the simplification **of linearization** (although some considerations of non-linearities are reintroduced at the end of chapter 4). But subject to this simplification, and under the basic conditions that the analysis in chapter 1 make explicit, the model is the right model to think about the allocation of resources in the world of private information we are considering.

2 A presentation of the model

Let us now introduce the basic model considered in this monograph, in particular in chapters 2 and 3. This model is a variant of the model of taxation studied by Diamond–Mirrlees in their pioneering study (1971); it is indeed known as **the Diamond–Mirrlees model**. It can be described as follows.

The economy has n private commodities indexed by $l = 1 \ldots n$ and one public good. Consumers indexed by $i = 1 \ldots m$ have preferences represented by a strictly quasi-concave utility function $U_i(z_i, y_0')$, where z_i is the vector of net trades of consumer i and y_0' the level of public good available in the society.

The economy consists of two firms with standard convex technologies: the first one is a private firm that has the standard competitive profit-maximizing behaviour, the second one is a semi-public firm that produces the specified level of public good y_0' together with the input vector $-y_0''$ (in other words the net output vector is y_0'') that is chosen in order to minimize the production cost of y_0' units of the public good.

The central feature of the model is that consumers, on the one hand, and producers, on the other, are faced with two different price systems. The first one, the **consumption price system**, is denoted π, the second one, the **production price system**, is denoted p. Both are vectors of $\mathscr{R}^n_+$. The difference $T = \pi - p$ is a vector of $\mathscr{R}^n$, called the tax vector or sometimes the vector of specific taxes.

Let us call $d_i(\pi, y_0')$ the excess demand of consumer i when the public good level is y_0' and the consumption price vector is π; let us call $\eta_1(p)$ and $\eta_0(y_0', p)$, respectively, the profit-maximizing and cost-minimizing supply function of the private and semi-public firm.

Formally:

$d_i(\pi, y_0')$ is the solution of the programme

$$\max U_i(z_i, y_0') | \pi . z_i \leq 0$$

$\eta_1(p)$ is the solution (also assumed to be unique) of

$$\max p.y_1'', y_1'' \in Y_1 \subset \mathscr{R}^n$$