Essays in Normative Macroeconomics

Thesis submitted for the Degree of DPhil

Charles Brendon Exeter College

Michaelmas Term 2011

Abstract

This thesis is divided into two main parts. The first provides a novel analysis of dynamic optimal taxation under the assumption that individuals in an economy have 'hidden' idiosyncratic productivity levels. Specifically, it shows how to derive a *complete* set of optimality conditions characterising the solution to a problem of this kind. The method relies on constructing perturbations to the consumption-output allocations of agents in a manner that preserves all relevant incentive compatibility restrictions. We are able to use it to generalise the 'inverse Euler condition' to cases in which preferences are non-separable between consumption and labour supply, and to prove a number of novel results about optimal income and savings tax wedges.

The second main part investigates a more general problem. When policymakers are constrained in their present choices by expectations of future outcomes a well-known time-inconsistency problem hinders optimal decision-making: the preferences of policymakers who exist at different points in time are not in agreement with one another, because of differences in the constraints faced by each. We present a new approach to determining policy in this setting, based on asking: What policy would be chosen by a decisionmaker who did not know the time period in which their choice was to be implemented? This is akin to designing institutions from behind a Rawlsian 'veil of ignorance'. The theory is used to obtain qualitative policy prescriptions across a number of environments; these policies have several appealing properties that we outline.

Contents

Part I

Introduction

The common theme that runs through this thesis is a focus on the dynamic aspects of optimal policy choice. More specifically, it is a focus on the *problems* for macroeconomic policymaking that are particular to the dynamic setting of certain models: how these problems can be isolated, and how best to treat them.

When choosing optimal policy interventions in dynamic macroeconomic models a time inconsistency problem will almost always feature. Indeed, so commonplace is this problem that it is now quite rare for normative papers even to draw explicit attention to it. Whenever a constraint set depends upon expectations across future allocations, optimal choice among these allocations from the perspective of an initial period will not in general also be best when viewed from the period in which the allocations themselves are observed. It is this problem in particular that sets time apart as more than 'just another dimension' by which to index the objects of policymaking.

One of the most important sources of expectational constraints is incentive compatibility. In many circumstances economic agents have private information, take unobservable actions, or simply have 'outside options' that require a policymaker to provide them with a minimum level of expected lifetime utility – so as to prevent deviation from a desired action profile. In general it will be most efficient to spread the provision of this utility through time, even though this may mean boosting an agent's welfare long after the *need* to incentivise them has passed. Where possible, it may also be desirable to change particular allocations in a way that alters prior expectations of the benefits from certain deviation strategies – again, even though the direct need to prevent this deviation has gone.

Part II of this thesis focuses on one particular model in which this logic applies: a dynamic Mirrleesian income tax model, of the type championed by the burgeoning New Dynamic Public Finance literature. The chief purpose of this Part is to highlight – and explore the implications of – a new analytical method for solving these models, under the assumption that optimal policy is to be evaluated from the perspective of an initial time period. This method relies on obtaining, through appropriately chosen marginal perturbations, a *complete* set of conditions sufficient to characterise an optimal allocation – obviating the need to apply dynamic programming techniques. The approach is readily generalisable, and is interesting in particular for the way it separates a large set of 'intratemporal' optimality conditions from a relatively small number of 'intertemporal' conditions. Analogous to a requirement in the analysis of optimal consumer choice that the marginal disutility of labour supply must always equal the real wage multiplied by the marginal utility of consumption, these 'intratemporal' conditions do not depend at all upon the dynamic setting of the model: they must equally apply in 'static' optimal income tax analyses.

The policy implications of the optimality conditions that we derive are very broad in scope. For instance, we show that optimal effective 'within-period' labour income taxes are always non-negative under an optimal allocation, and that it will almost always be optimal for policy to deter savings at the margin (in a specific sense that is made clear) – though when consumption and labour supply are Edgeworth complements this need not be the case. We are also able to draw particular attention to the extra richness of optimal policy when productivity disturbances are non-iid. In this situation there are two, rather than one, intertemporal optimality conditions, whilst the number of intratemporal requirements is reduced by one correspondingly. The extra dynamic requirement derives from the capacity of the policymaker to exploit differences in the productivity measures of truth-telling agents and those who would potentially mimic them, and illustrates clearly the 'theory of the second best': where possible it is better to introduce extra distortions to market allocations so long as these reduce the size of other inefficiencies at the margin.

To the extent that time-inconsistency problems do apply, the work presented in Part II allows them to be identified exclusively with satisfaction of the *inter*temporal optimality conditions. It is these conditions alone that are particular to the dynamic setting, and the novel method that is outlined in Part II allows the majority of (relatively innocuous) intratemporal requirements to be viewed in isolation from them. Moreover, we show that under many plausible preference assumptions a direct implication of one of the dynamic conditions (the so-called 'inverse Euler equation', which we are able to generalise substantially) is that 'immiseration' will take place almost surely for all agents in the economy as time progresses, in the sense that their consumption will almost surely converge to zero under the optimal allocation scheme. This startling result – already well known when preferences are separable between consumption and labour supply – must be interpreted at least as evidence against the *implementability* of any such scheme: it is scarcely credible that agents in an economy would willingly accept servitude simply because it was best that they should do so from the perspective of a policymaker in the distant past. And this point is quite distinct from the many normative issues that so unequal an intergenerational allocation highlights: why *should* the initial period's policymaker have the right to dictate outcomes in perpetuity? More generally, how should the divergent preferences of policmakers operating at different points in time be more appropriately reconciled?

It is to questions such as these that Part III of the thesis turns. There are in fact strong *practical* grounds for addressing them. Suppose, for instance, that a central bank were delegated to pursue a monetary policy strategy that was deemed optimal from the perspective of the specific time period in which it was granted independence. Then the best inflation rate for it to target would generally differ depending entirely on the time elapsed since independence – holding constant the current and past state of the economy. But no central bank does, in practice, follow a time-varying strategy such as this, and no prominent policymaker has advocated one. In short, there is a disconnect between normative theory and practice, with a practical bias in favour of *time-invariant* policy approaches. This is true not just of monetary policy: there seem no obvious examples in which taxation or public expenditure strategies have been deliberately formulated to change through time in a way that would have been to the particular advantage of policymakers operating in the period in which they were first devised.

To date the work of Michael Woodford has done most to address this disconnect. Focusing chiefly on New Keynesian models of optimal monetary and fiscal policy, Woodford has championed a 'timeless perspective' approach to policy choice. Under this method, all outcomes must be chosen according to the preferences of a hypothetical policymaker alive in the infinitely distant past. Put differently, policymakers should *immediately* implement the history-dependent allocations that would otherwise obtain only at the limit as time goes to infinity under an optimal plan from the perspective of an initial period.

But although Woodford's method succeeds in achieving time-invariance, it is not clear it provides the best possible resolution. Two important practical problems with it are highlighted in Part III. First, in some quite simple models the timeless perspective policy is unambiguously dominated within the set of possible time-invariant strategies by alternative approaches. Second, in many models with infinite-horizon incentive-compatibility constraints (including the dynamic tax models that feature in Part II) the timeless perspective strategy is not well defined. More broadly, it is not clear why 'Do what would have been best from the perspective of the distant past.' should be considered a particularly appealing normative criterion.

Instead, the work in Part III asks what actions would be taken by policymakers who were denied all knowledge of the time period for which they were choosing. We argue that this is a more natural perspective to take. In the general setup that we consider, time inconsistency derives entirely from the fact that a policymaker in period t knows that expectations for periods $t-1$ and earlier can no longer be affected by the current choice (whereas these expectations are still to be determined in the event that earlier policymakers are setting plans for period t). Apart from these differences in the set of *constraints* that they face, policymakers in different time periods are identical. John Rawls's famous 'veil of ignorance' construct has provided an influential way for political theorists to devise normative principles for institutional design when the would-be designers differ in their particular circumstances. Specifically, Rawls advocates the design that would be selected by all agents in the event that they were denied knowledge of their own circumstances – being forced to choose as if from behind the veil of ignorance. This principle has clear and direct applicability in the context of the time-inconsistency problems that we study: for these problems, policymakers differ only according to the time period in which they assess outcomes. Forcing them to make choice without knowing the period for which they are choosing will deliver coherence across their preference rankings.

Exactly how to deny knowledge of time is a complicated problem, for reasons explored in great detail in Part III. Broadly, the difficulty lies in choosing the *minimal* set of information that it is sufficient to deny to the policymaker, making sure time cannot be known or inferred whilst still permitting rich *history* dependence in policy. The notation inevitably becomes rather involved, but we are able to present a well-defined choice procedure for determining the allocations that we dub 'best from behind a veil of ignorance'. Importantly, this procedure will select unambiguously better allocations to the timeless perspective approach in models where a comparison can clearly be made (that is, purely deterministic models with no endogenous state variables), and will be well defined across a far broader class of problems – allowing infinite-horizon incentive-compatibility constraints to be incorporated relatively straightforwardly.

The policy method that we specify in this way has subtle but important consequences. If it is unclear *when* a particular policy will operate, the dynamic weighting of distinct marginal effects changes, with the effects of changes to prior expectations ultimately given slightly less weight than would be the case under optimal policy from the perspective of an initial time period. The intuition behind this is explained in more detail below: it rests crucially on understanding the time preference structure of a policymaker who chooses outcomes that are best from an initial period. The most significant *implication* of the veil of ignorance approach is that many of the dynamic optimality conditions that induce non-stationarity in target variables when solving for the best initial-period policy – particularly in models with incentive compatibility constraints – are instead now consistent with *stationarity*. A notable example is the generalised 'inverse Euler condition' derived in Part II. This condition was itself behind one of the most extreme examples of a time-inconsistent allocation, with almost all agents consuming nothing at the limit as time progressed under the best dynamic tax policy from the perspective of an initial period. The veil of ignorance allocation instead involves simple mean-reversion in consumption and output for all agents, with a post-tax income distribution that is stable through time. This is a policy that still exploits the benefits from spreading incentives through time, but does so without inducing severe disparities in the welfare of different generations.

Moreover, we can show that all of the intratemporal optimality conditions derived for the optimal tax model in Part II will still apply at a veil of ignorance optimum, reinforcing the value of having separated these from the more contentious dynamic restrictions. A similar result applies for a simple model of redistribution when agents receive stochastic incomes and are imperfectly committed to the scheme: relative consumption levels between agents with given 'Pareto weights' *within* a particular period are exactly as they would be under an optimal policy from the perspective of an initial time period. But these consumption levels (and Pareto weights) evolve differently over time.

Ultimately there can never be a perfect way to resolve time inconsistency problems. The 'veil of ignorance' method presented in Part III of this thesis is not any more 'correct' than assuming that a policymaker in an initial time period can (or should) impose his or her will at all future horizons. But the work in Part II reiterates the discomfiting implications that this latter approach can have. If optimal policymaking criteria are themselves to be judged by their practical implications, the veil of ignorance approach may well have much to commend it.

Part II

Applying perturbation analysis to dynamic optimal tax problems

1 Introduction

There is a growing interest among macroeconomists in dynamic optimal policy problems in the presence of asymmetric information. One such class of problems that has received particular attention is that of multi-period optimal tax analyses, based on the seminal works of Mirrlees (1971) and Diamond and Mirrlees (1978). Yet the complexity of the models in which this analysis is conducted has led to relatively few general analytical results emerging, of the kind that might confidently inform policy discussions.[1] Considerable progress has certainly been made under special assumptions regarding utility functions and skill distributions,[2] but in what ways the associated results generalise remains an open question. Indeed, the most clear (and most celebrated) analytical statement that *has* emerged – the so-called 'inverse Euler condition' – is itself particular to the quite strict requirement that consumption and labour supply should be separable in all agents' preferences. In short, there is much theoretical work still to be done.

The aim of this Part is to contribute to that theoretical project. Working under the assumption that the 'first-order approach' is valid – so that the set of incentive compatibility constraints that binds at the optimum is known – we set out a novel perturbation method that is capable of providing a complete characterisation of that optimum. That is, we are able to obtain a set of distinct necessary

[1] In the words of one prominent recent survey (Mankiw, Weinzierl and Yagan (2009)): "The theory of optimal taxation has yet to deliver clear guidance on a general system of history-dependent, coordinated labor and capital taxation ... Most of the recommendations of dynamic optimal tax theory are recent and complex."

[2] Important recent contributions in this regard are Farhi and Werning (2010) and Golosov, Troshkin and Tsyvinski (2011).

optimality conditions exactly equal in number to the degrees of freedom available to the policymaker. In itself this result holds out the promise of substantially simplifying the numerical calculation of optimal tax schedules, which reduces to a 'mechanical' question of solving a given – albeit often very large – system of simultaneous equations (with no need to use dynamic programming techniques). Perhaps more significantly, the conditions that we derive imply several important general results regarding the character of dynamic optimal tax schedules.

First, and of substantial theoretical interest, we are able to generalise the inverse Euler condition to situations in which an agent's within-period consumption and labour levels are non-separable in utility. Second, using this result we are able to reach the general conclusion that optimal taxes should always deter savings (in a well-defined sense) when consumption and labour supply are either substitutes or separable from one another in preferences, but that this need not be true in the event that they are complements. Third, and again using our generalisation of the inverse Euler condition, we show that the long-run 'immiseration' results that are known to characterise dynamic Mirrlees economies with infinitely-lived dynasties under preference separability again generalise to the case of substitutes, but not of complements. The previous two results seem closely related, and shed some light on the precise dynamics responsible for immiseration.

Fourth, we show that the set of 'intratemporal' optimality conditions characterising allocations when skill distributions are iid is identical to the set of conditions that must hold in a static optimal income tax model, providing an important mapping between the traditional 'public economics' and more recent 'mechanism design' literatures.[3] But (fifth) when skills are Markov in a more general sense there is a reduction by one in the number of intratemporal optimality conditions – supplanted by an additional *inter*temporal condition, capturing the capacity of the policymaker to spread through time the distortions required to prevent more productive agents from mimicking. This condition will generally imply higher marginal tax rates as time progresses for those whose productivity has been low.

[3] The distinction is drawn by Diamond and Saez (2011).

Sixth, we show that it is never optimal in any time period to subsidise labour supply at the margin. Seventh (and finally) we show that effective marginal labour income tax rates will always be zero at the top of the skill distribution, in the event that this distribution has an upper support.

The key argument that lies behind all of these results is that if the first-order approach is valid then it is always possible to construct a set of perturbations to optimal (equilibrium) allocations such that *local* incentive compatability constraints will continue to bind. That is, if we know that agent A is just envied by agent B in equilibrium (so that truth-telling is only weakly preferred to mimicking by the latter), we can construct simultaneous changes to the consumption and income levels of each such that the resulting increase in agent B's utility from truth-telling is exactly the same as any increase in the utility he or she could obtain by mimicking agent A. If an allocation is optimal, such perturbations cannot be used to generate surplus resources, provided they respect all *prior* incentive compatibility constraints.

Specifically, this approach requires that one should define perturbations to the optimal allocation that simultaneously satisfy three conditions: local incentive compatibility, reversibility, and welfare-neutrality. The first of these is particular to dynamic screening models: a perturbation to outcomes that changes the incentives for truthful reporting at the same time as it changes allocations will not generally be of use for our purposes, due to the discrete shifts in consumption and output patterns that would follow as agents change their reports. We are interested, rather, in studying perturbations to *allocations* whilst holding constant agents' *type reports* (exploiting the revelation principle to focus on a mechanism whereby agents report their idiosyncratic productivities directly).[4]

The second of the conditions is necessary if optimality conditions are to be stated with equality. It demands that if we can increase the consumption and output allocations of agents along some vector Δ at the margin, then we can also increase them along the vector $-\Delta$. In a simple consumption-savings

[4]The 'static' optimal income tax literature also makes use of perturbation analysis, but without exploiting direct revelation mechanisms: rather, the focus is directly on changes to the tax *schedule* subject to which all individuals choose. See, in particular, Roberts (2000) and Saez (2001).

problem, this is the equivalent of noting that we must not be at a corner solution if we are to state the consumption Euler equation with equality.

The final requirement is that the perturbations should be welfare-neutral from the perspective of the policymaker in the initial time period. This is useful, since it means we can focus simply on whether any given perturbation raises surplus resources in assessing whether it is to the advantage of the policymaker. Satisfying these three requirements for a broad class of perturbations – far broader than the intertemporal utility reallocations already applied in the literature when deriving the inverse Euler condition – is a non-trivial challenge, and establishing a general procedure for doing so forms the heart of the analysis in what follows.

1.1 Literature review

To date, two closely related methodological approaches to solving dynamic problems under asymmetric information have emerged in the macroeconomics literature. The first, and most widely-used, follows the foundational work on dynamic games by Abreu, Pearce and Stachetti (1990), considering directly the planner's problem of maximising a given objective criterion subject to a series of lifetime utility constraints that must hold in each time period in equilibrium (preventing any incentive for agents to mis-report their private information). Examples include Atkeson and Lucas (1992), investigating consumption allocations across agents subject to idiosyncratic taste shocks; Kocherlakota (1996), looking at consumption risk sharing when incomes are stochastic; Golosov, Troshkin and Tsyvinski (2011) in a dynamic tax setting; and numerous other papers besides. An important feature of these approaches is the reformulation of the policymaker's problem into an equivalent recursive choice across current outcomes and a vector of discounted utility promises – the latter summarising the dynamic incentives that are being provided to ensure truthful reporting.

An important refinement to this method – particularly in the context of the present paper – has been provieded by Kapička (2010) (extending the general work of Pavan, Segal and Toikka (2011)), who illustrated the potential of the 'first-order approach' to reduce the state-space required in dynamic Mirrlees models – particularly in the (realistic) event that agents' productivities evolve according to non-iid processes. Specifically, Kapička demonstrates that one requires just two variables to summarise the policymaker's past promises to an individual with a given history of productivity draws: a promised lifetime utility, and a value expressing how this utility changes at the margin as the agent's type changes. This method substantially eases the computational burden associated with calculating optimal allocations by the 'primal' (promised utilities) recursive technique, relative to existing methods valid under non-iid assumptions – notably that of Fernandes and Phelan (2000). It has been adopted fruitfully by Farhi and Werning (2010) among others.

The second general approach, referred to as the 'dual' method by Messner, Pavoni and Sleet (2011), follows Marcet and Marimon (1998) in exploiting the evolution of *costates* associated with lifetime utility constraints, in order to augment the policymaker's objective criterion in a manner that ensures incentive compatibility constraints are always satisfied. The problem is again set in a recursive form, but with no explicit choice over a set of future utilities; instead the Pareto weights that are placed on distinct agents' utilities in the policy objective are increased exactly as necessary to ensure the resulting optimisation satisfies incentive compatibility. This method has recently been applied to optimal dynamic tax policy by Sleet and Yeltekin (2010b), and relatedly by Mele (2011) to hidden action problems. The latter authors have also provided a general analysis applying the earlier theory to settings with private information (see Sleet and Yeltekin (2010a)), as has recent additional work by Marcet and Marimon (2011).

Both of these methods arrive at solutions to the underlying problem through functional iteration on a Bellman-type operator. Whilst this has the advantage of quite widespread applicability, it necessitates numerical methods that may prevent the essental analytical character of the solution from being completely clear. Rather than follow these papers in pursuing a variant upon the dynamic programming literature, here we instead develop a method more closely related to the calculus of variations. That is, we assume that an optimum has been found, and ask what properties that optimum would have to satisfy. This logic has already been applied by Kocherlakota (2005) and Golosov, Tsyvinski and Werning (2006), among others, to obtain one particular necessary optimality condition in a dynamic Mirrleesian economy for which preferences between leisure and consumption are separable: here it has been show that an 'inverse Euler condition' must hold, linking the marginal cost of providing consumption utility to a consumer in one time period to the expected value of the same marginal cost across distinct realisations for that consumer's idiosyncratic productivity level in the next period. The marginal cost of providing consumption utility is the inverse of the marginal utility of consumption. The basic idea is that if an allocation is optimal the policymaker cannot transfer through time the provision of a unit of utility to a consumer with a particular productivity history and raise a resource surplus.

Since the work of Thomas and Worrall (1990) (in the context of a repeated moral hazard model) the long-run implications of intertemporal optimality conditions of this form have been an important focus of study, particularly in the event that they define a bounded martingale sequence – which is generally the case when the real rate of interest is equal to consumers' and the policymaker's discount factor.[5] In this case long-run 'immiseration' results will generally follow by applying matingale convergence theorems, in the sense that the marginal utility of consumption for all agents will almost surely be unboundedly large as time progresses (its inverse will almost surely converge on zero). But these results rely heavily on the assumption of separability between consumption and labour supply. In the event that this does not hold, the policymaker may find that an incentive-compatible perturbation requires changes to the labour supply of agents simultaneously to consumption changes, with the implication that the inverse of the marginal utility of consumption is no longer the marginal cost of (incentive-compatible) utility provision. We show that this is indeed the case in our optimal tax model, and

[5] Papers by Phelan (2006) and Farhi and Werning (2007) consider the implications of the social discount factor differing from individuals', showing that the inverse Euler condition ceases to be a martingale in this event, so long as the real interest rate remains equal to the inverse of the *household* discount factor.

explore the dynamic implications that follow.

2 Model setup

The basic framework that we use essentially follows the recent textbook treatment of Kocherlakota (2011), except that we allow for a general specification of preferences from the outset. An economy is populated by a large number of agents, modeled as a continuum with each agent indexed by a position on the unit interval. Each agent is the current manifestation of an infinitely-lived dynasty, and gains utility from that dynasty's expected consumption and leisure from the current period into the infinite future. Labour is the only factor of production and there are no firms – so agents can be thought of as directly choosing the level of output that they produce each period via their labour supply decision. Their preferences over output and consumption profiles from time t onwards are described by the function U_t:

$$U_t = E_t \sum_{s=0}^{\infty} \beta^s u\left(c_{t+s}, y_{t+s}; \theta_{t+s}\right) \tag{2.1}$$

where $u : \mathbb{R}^3_+ \rightarrow \mathbb{R}$. c_{t+s} and y_{t+s} are, respectively, the agent's consumption and output levels in period $t+s$, $\beta \in (0,1)$ is the dynasty's time preference parameter, and θ_{t+s} is an idiosyncratic productivity parameter that allows one to map from a level of output to a quantity of labour supply. The productivity parameter belongs to a set $\Theta \subset \mathbb{R}$, which is time- and history-invariant.[6] For the entirety of this paper we work under the assumption that Θ contains a finite number of elements N, which turns out to provide the most straightforward setting in which to present the main arguments. We generalise to the (more conventional) assumption that Θ is an interval of the real line in a companion paper (see Brendon (2011)). We also assume an infinite horizon, though none of the optimality conditions that

[6] The analysis is made simpler by assuming that Θ itself does not depend on past draws. The probability of any one element of Θ being drawn after a given history can always be made arbitrarily small, so this does not seem a particularly restrictive assumption.

we derive is dependent on this perspective.[7]

Expectations are taken across all stochastic variables relevant to the equilibrium evolution of the agent's utility (ultimately, drawings from Θ at each future horizon). We analyse the model as if nature is responsible at the start of time for drawing a distinct element for each dynasty from the infinite-dimensional product space Θ^{∞}, say θ^{∞}, according to some probability measure on Θ^{∞}, π_{Θ}. These draws are iid across dynasties. At the start of each time period agents are informed of their within-period productivity, so that at time t they are aware of their complete history of draws to date, $\theta^t \in \Theta^t$, where θ^t is a t-length truncation of θ^{∞}. This purely idiosyncratic information is private knowledge to the agent, so policymakers must provide sufficient incentives to prevent mimicking in any tax system that is established.

To keep the problem regular we assume that the utility function is twice continuously differentiable in all of its arguments, with $u_c > 0$, $u_y < 0$, and $u_\theta > 0$, and that the partial Hessian $\begin{bmatrix} u_{cc} & u_{cy} \\ u_{cy} & u_{yy} \end{bmatrix}$ is negative definite for any given θ. We additionally impose Inada conditions: $\lim_{c\to\infty} u_c(c, y; \theta) = 0$ and $\lim_{c\to 0} u_c(c, y; \theta) = \infty$ for all non-zero, finite (y, θ) pairs, and $\lim_{y\to\infty} u_y(c, y; \theta) = -\infty$ and $\lim_{y\to 0} u_y(c, y; \theta) = 0$ for all non-zero, finite (c, θ) pairs. These conditions will ensure an interior solution obtains at all finite horizons. To maintain the interpretation of the model as an optimal tax problem with unobservable labour supply we impose that marginal changes to θ will reduce the marginal disutility associated with a unit of extra output when consumption and utility (and thus, implicitly, labour) are jointly held constant. This can be shown to imply:

$$u_{y\theta} - u_{yy}\frac{u_\theta}{u_y} > 0 \tag{2.2}$$

Similarly, if consumption and utility are jointly held constant as θ is changed then labour supply must implicitly also be being held fixed – and thus the marginal utility of consumption should likewise

[7] An interesting feature of our approach is that it provides a novel representation of the optimality requirements even in a 'static' optimal income tax model.

be constant. This is quite easily shown to imply the following:

$$u_{c\theta} - u_{cy}\frac{u_\theta}{u_y} = 0 \tag{2.3}$$

Finally, a variant upon the Spence-Mirrlees single-crossing condition is imposed, to ensure higher realisations of θ can naturally be associated with higher productivity:

$$u\left(c'', y''; \theta''\right) - u\left(c', y'; \theta''\right) > u\left(c'', y''; \theta'\right) - u\left(c', y'; \theta'\right) \tag{2.4}$$

whenever $c'' > c'$, $y'' > y'$ and $\theta'' > \theta'$.

Note that this condition is slightly stronger than could be obtained simply by differentiating the expression for the slope of a within-period indifference curve in output-consumption space ($\frac{dc}{dy} = -\frac{u_y}{u_c}$); although (2.4) *implies* that this indifference curve should be flattening in θ (as seen by assuming one of the agents is indifferent between the two bundles), it also implies certain properties are associated with utility changes between bundles across which neither agent is indifferent, and we exploit these properties to some extent in what follows (notably when stating sufficient conditions for the 'first-order approach' to be consistent with global incentive compatibility). Occasionally it is useful also to state the condition in terms of marginal rates of substitution: if $\theta'' > \theta'$ then condition (2.4) implies for all (c, y) pairs:

$$-\frac{u_y\left(c, y; \theta''\right)}{u_c\left(c, y; \theta''\right)} < -\frac{u_y\left(c, y; \theta'\right)}{u_c\left(c, y; \theta'\right)} \tag{2.5}$$

This follows directly from the fact that indifference curves in consumption-output space must be 'flattening' as θ increases, provided (2.4) holds.

The policymaker's role is to choose, at the start of the first time period, effective tax schedules for all future periods that will link an individual's consumption to their output, conditional on their history of actions to date. The purpose of this choice is to maximise some social welfare function, defined across the set of possible equilibrium allocations. Individuals can be thought of as devising history-contingent action profiles to implement in each future time period, given the mechanism with which the policymaker presents them. Since the revelation principle will apply in this setting,[8] we may restrict policy choice to direct revelation mechanisms that deliver consumption and output bundles to individuals as functions of their current and past productivity reports – deferring a consideration of decentralisation schemes for later work. In treating consumption as a choice variable of the policymaker in this way we are implicitly assuming that there are no opportunities for the individuals to engage in 'hidden' saving – so that the policymaker could always behave as if taxing savings at a 100 per cent marginal rate, if this were necessary to prevent 'unwanted' consumption deferral.[9] We generally denote by $\widehat{\theta}_t^i : \Theta^t \to \Theta$ individual i's report at time t as a function of their actual productivity (where this function is measurable with respect to θ^t), by $\widehat{\theta}^{i,t} : \Theta^t \to \Theta^t$ the history of all such reports up to time t, and by $\widehat{\theta}^{i,\infty} : \Theta^\infty \to \Theta^\infty$ a complete sequence of reports. We occasionally refer to $\widehat{\theta}^{i,t}(\cdot)$ as the t-truncation of $\widehat{\theta}^{i,\infty}(\cdot)$.

For the remainder of the paper we follow the majority of the literature and assume a utilitarian policy criterion, assessed from the perspective of the initial time period. This criterion has the advantage in a dynamic context of being the only objective that satisfies 'welfarism' at every horizon that is also time-consistent. That is to say, social welfare is assessed in each period as a function of individual lifetime utilities alone, and if two candidate policies deliver exactly the same outcomes between periods 1 and t then the relative preference of the policymaker between those two paths will be the same at time 1 as at time t. Whilst no claim is made that these normative features should be elevated above all others, they do arguably allow for the simplest treatment of the dynamic tax questions that are of chief interest to us.

[8] We seek a Bayes-Nash equilibrium of the game played between the policymaker and all individuals whose types may be drawn from Θ^∞. The revelation principle states that any such equilibrium can be supported by a direct revelation mechanism.

[9] Da Costa and Werning (2002) and Golosov and Tsyvinski (2006) consider economies with hidden savings opportunities; these substantially reduce the options available to the policymaker.

The policymaker's primal choice problem is, therefore:

$$\max_{\{c_t(\theta^\infty), y_t(\theta^\infty)\}_{t=1}^{\infty}} \int_{\Theta^\infty} \sum_{t=1}^{\infty} \beta^{t-1} u\left(c_t\left(\theta^\infty\right), y_t\left(\theta^\infty\right); \theta_t\right) d\pi_\Theta\left(\theta^\infty\right) \tag{2.6}$$

subject to $c_t(\theta^\infty)$ and $y_t(\theta^\infty)$ being measurable with respect to θ^t, together with the incentive compatibility constraint:

$$\begin{aligned} & \int_{\Theta^\infty} \sum_{s=0}^{\infty} \beta^s u\left(c_{t+s}\left(\theta^\infty\right), y_{t+s}\left(\theta^\infty\right); \theta_{t+s}\right) d\pi_\Theta\left(\theta^\infty | \theta^t\right) \\ \geq & \int_{\Theta^\infty} \sum_{s=0}^{\infty} \beta^s u\left(c_{t+s}\left(\widehat{\theta}^\infty\left(\theta^\infty\right)\right), y_{t+s}\left(\widehat{\theta}^\infty\left(\theta^\infty\right)\right); \theta_{t+s}\right) d\pi_\Theta\left(\theta^\infty | \theta^t\right) \end{aligned} \tag{2.7}$$

which must hold for all t, all θ^t, and all functions $\widehat{\theta}^\infty : \Theta^\infty \to \Theta^\infty$ whose s-truncations $\widehat{\theta}^s(\cdot)$ are measurable with respect to θ^s for all $s \geq 1$; and finally the resource constraint:

$$\int_{\Theta^\infty} \left[c_t\left(\theta^\infty\right) - y_t\left(\theta^\infty\right)\right] d\pi_\Theta\left(\theta^\infty\right) + A_{t+1} = R_t A_t \tag{2.8}$$

where A_t is the quantity of real bonds that the policymaker purchases for time t, each paying R_t units of real income in that period. The initial value $R_1 A_1$ is given. Dynamic solvency requires that $\lim_{s\to\infty}\left[\left(\prod_{r=1}^{s} R_{t+r}^{-1}\right) A_{t+s}\right] = 0$.[10]

3 Full information benchmark

In a manner equivalent to Kapička (2010) and Broer, Kapička and Klein (2011), we will ultimately focus our attention on a relaxed version of the incentive compatibility constraint, arguing (in the context of a discrete number of types in Θ) that it is sufficient to impose a binding restriction to prevent agents with histories $\left(\theta^{t-1}, \theta_t\right)$ mimicking those with histories $\left(\theta^{t-1}, \theta_t'\right)$, where $\theta_t' = \max\left\{\theta \in \Theta : \theta < \theta_t\right\}$.

[10] In what follows it is often convenient to suppress the explicit dependence of c_t and y_t upon θ^∞; we also occasionally index these functions with individual-specific superscripts where this is most natural.

The basic reason for our making this assumption – that envy is always directed 'downwards' from one type to the next in equilibrium – is familiar from the analysis of static optimal tax models, and was articulated most clearly by Dasgupta (1982). To understand why it is likely to hold, it is useful to start by considering the character of optimal policy when the idiosyncratic productivity draws are common knowledge.

If the policymaker is aware of agents' types each period the incentive compatibility constraint (2.7) can be neglected, with lump-sum taxation used to implement a first-best. We summarise four important properties of this first-best in the following list. The proofs of each statement are trivial, and hence omitted, with the exception of the fourth, which is provided in the appendix.

1. In the full information benchmark the optimal allocations $c_t(\theta^\infty)$ and $y_t(\theta^\infty)$ are measurable with respect to θ_t.

2. The following conditions hold for all $t \geq 1$ and all $i \in [0,1]$:

$$u_c\left(c_t^i, y_t^i; \theta_t^i\right) = -u_y\left(c_t^i, y_t^i; \theta_t^i\right) \tag{3.1}$$

$$u_c\left(c_t^i, y_t^i; \theta_t^i\right) = \beta R_{t+1} \sum_{\theta_{t+1}^i \in \Theta} u_c\left(c_{t+1}^i, y_{t+1}^i; \theta_{t+1}^i\right) \pi_\Theta\left(\theta_{t+1}^i | \theta_t^i\right) \tag{3.2}$$

3. The following condition holds for all $t \geq 1$ and all agents $i, j \in [0,1]$:

$$u_c\left(c_t^i, y_t^i; \theta_t^i\right) = u_c\left(c_t^j, y_t^j; \theta_t^j\right) \tag{3.3}$$

   In the event that consumption and labour are additively separable in the utility function we will additionally have $u_{cy} = u_{c\theta} = 0$, and this condition then implies equalised consumption across all agents (since $u_{cc} < 0$).

4. $\theta_t^i > \theta_t^j$ implies $u\left(c_t^i, y_t^i; \theta_t^i\right) < u\left(c_t^j, y_t^j; \theta_t^j\right)$, so long as leisure is a normal good (at autarky

prices).

Summarising the main lessons of these four statements in turn, we know from the first that there is no incentive for the policymaker to introduce any form of history dependence in outcomes. The fact that a particular individual has been very productive in the past makes no difference to their optimal current consumption-output bundle, independently of the contemporary productivity draw θ_t. In this sense the first-best solution offers no scope for agents to claim credit for past accomplishments. The second statement implies that the optimal solution for a utilitarian policymaker involves zero marginal distortions on savings and labour supply, whilst the third points to equalised marginal consumption utility (and, thus, output disutility) across agents each period. Since agents who are more productive have, by definition, a higher marginal product for a *given* quantity of labour they will generally be required to work longer hours at the optimum. This is the logic behind the fourth condition – that utility is decreasing in type so long as leisure is a normal good.[11] This last result is key to understanding which incentive compatibility constraints will bind at the optimum: together with the fact that there is no history dependence in outcomes at the first-best, it strongly implies higher-type agents would mimic their lower-type peers if they had the capacity to do so – that is, in the event that the policymaker could only verify agents' output levels, and not their types.

4 The first-order approach to incentive compatibility

We now move to the constrained problem, in which the policymaker is forced to abide by incentive compatibility constraints – and hence will be prevented from providing higher-productivity types with a lower level of utility than their (lower-productivity) peers. As mentioned above, we retain a focus on the case in which Θ contains a discrete, finite number of elements. To apply our perturbation method, we first need to be clearer on the set of constraints that will bind at the optimum.

[11] It is, indeed, a case of 'From each according to his means, to each according to his needs.'

For all periods $t \geq 1$, define $\widehat{\theta}_{m,t}^{\infty} : \Theta^{\infty} \rightarrow \Theta^{\infty}$ as the reporting strategy associated with truth-telling in all periods up to t, at which point the agent mimics a type one lower and follows an optimal reporting strategy thereafter:

$$\widehat{\theta}_{m,t}^{\infty}\left(\theta'^{\infty}\right) = \left[\theta'_1, \theta'_2, ..., \theta'_{t-1}, \theta''_t, \theta''_{t+1}, ...\right]$$

where $\theta''_t = \max\left\{\theta \in \Theta : \theta < \theta'_t\right\}$ and $\left\{\theta''_{t+1}, \theta''_{t+2}, ...\right\}$ are then optimal choices conditional upon prior reports. So long as the type distribution is Markov, outcomes for an agent with a given reporting history will in fact be independent of the veracity of that reporting history – so we are free to focus exclusively on 'one-off' deviations from the truth, with $\left\{\theta''_{t+1}, \theta''_{t+2}, ...\right\} = \left\{\theta'_{t+1}, \theta'_{t+2}, ...\right\}$. If $\theta'_t = \min\left\{\theta \in \Theta\right\}$ then we normalise $\widehat{\theta}_{m,t}^{\infty}\left(\theta^{\infty}\right) = \theta^{\infty}$. If incentive compatabilty is said to be holding 'downwards', the following is true:

$$\begin{aligned} & \int_{\Theta^{\infty}} \sum_{s=0}^{\infty} \beta^s u\left(c_{t+s}\left(\theta^{\infty}\right), y_{t+s}\left(\theta^{\infty}\right); \theta_{t+s}\right) d\pi_{\Theta}\left(\theta^{\infty}|\theta^t\right) \\ \geq \; & \int_{\Theta^{\infty}} \sum_{s=0}^{\infty} \beta^s u\left(c_{t+s}\left(\widehat{\theta}_{m,t}^{\infty}\left(\theta^{\infty}\right)\right), y_{t+s}\left(\widehat{\theta}_{m,t}^{\infty}\left(\theta^{\infty}\right)\right); \theta_{t+s}\right) d\pi_{\Theta}\left(\theta^{\infty}|\theta^t\right) \end{aligned} \tag{4.1}$$

So the agent with history θ^t weakly prefers reporting θ_t truthfully to mimicking a type one lower, provided θ_t is not itself the smallest element in Θ. Again, for any Markovian productivity process it must be true that if (4.1) holds for agents whose past reports of $\widehat{\theta}^{t-1}$ were truthful, it must hold for *all* agents with past reports of $\widehat{\theta}^{t-1}$ and a true contemporary productivity draw equal to θ_t. We are interested in the conditions under which this restriction implies *global* incentive compatibility – that is, for an arbitrary reporting strategy at t, $\widehat{\theta}_{a,t}^{\infty} : \Theta^{\infty} \rightarrow \Theta^{\infty}$, defined by:

$$\widehat{\theta}_{a,t}^{\infty}\left(\theta'^{\infty}\right) = \left[\theta'_1, \theta'_2, ..., \theta'_{t-1}, \theta''_t, \theta''_{t+1}, ...\right]$$

for any $\theta''_t \in \Theta$, with $\left\{\theta''_{t+1}, \theta''_{t+2}, ...\right\}$ chosen optimally thereafter, we want to know when it will be the case that inequality (4.1) implies:

$$\begin{aligned} & \int_{\Theta^\infty} \sum_{s=0}^{\infty} \beta^s u\left(c_{t+s}\left(\theta^\infty\right), y_{t+s}\left(\theta^\infty\right); \theta_{t+s}\right) d\pi_\Theta\left(\theta^\infty | \theta^t\right) \\ \geq & \int_{\Theta^\infty} \sum_{s=0}^{\infty} \beta^s u\left(c_{t+s}\left(\widehat{\theta}_{a,t}^\infty\left(\theta^\infty\right)\right), y_{t+s}\left(\widehat{\theta}_{a,t}^\infty\left(\theta^\infty\right)\right); \theta_{t+s}\right) d\pi_\Theta\left(\theta^\infty | \theta^t\right) \end{aligned} \tag{4.2}$$

This problem lies at the heart of discussions on the applicability of the first-order approach in problems of this kind – an issue first considered by Mirrlees (1971), and studied in great depth in the context of dynamic models by Pavan, Segal and Toikka (2011). The first-order approach takes as its starting point the fact that under any incentive-compatible direct revelation mechanism no agent can induce an increase in their expected lifetime utility by changing their report. We can define the value function $W\left(\widehat{\theta}_t; \theta_t, \widehat{\theta}^{t-1}\right)$, with $W : \Theta \times \Theta \times \Theta^{t-1} \to \mathbb{R}$ specifying the maximum lifetime utility that could be expected for an agent whose past reports were $\widehat{\theta}^{t-1}$, whose current productivity is θ_t and whose current report is $\widehat{\theta}_t$. Then the approach notes that for a given $\left(\theta_t, \widehat{\theta}^{t-1}\right)$ pair this function must have a global maximum where $\widehat{\theta}_t = \theta_t$. Thus instead of choosing directly from among the (difficult to characterise) set of allocations for which condition (4.2) is explicitly asserted for all admissible functions $\widehat{\theta}_{a,t}^\infty$, one may instead choose simply from the set for which $W\left(\cdot; \theta_t, \widehat{\theta}^{t-1}\right)$ is known to have a stationary point at θ_t. In the case that a discrete number of types features in Θ (rather than Θ being a proper subset of the real line), it is not directly apparent what this implies: we cannot place a restriction on the derivative of W with respect to $\widehat{\theta}_t$ if there is no possibility of marginal changes to the agent's report. Yet we may invoke our earlier result that the first-best optimum involves decreasing utility in θ to apply a 'first-order' approach in which choice is from the set of allocations such that the condition:

$$W\left(\theta_t; \theta_t, \widehat{\theta}^{t-1}\right) \geq W\left(\theta_t'; \theta_t, \widehat{\theta}^{t-1}\right) \tag{4.3}$$

is imposed for $\theta_t' = \max\left\{\theta \in \Theta : \theta < \theta_t\right\}$.[12] That is, consistent with the familiar logic of the Mirrlees

[12] No restriction is imposed in the event that $\theta_t = \min\left\{\theta \in \Theta\right\}$.

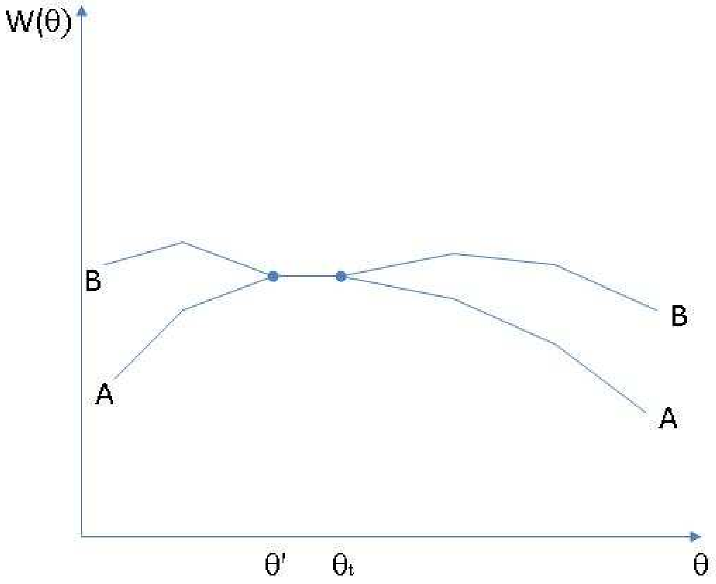

Figure 4.1: Local incentive compatibility need not imply global.

model, incentive compatibility must be imposed 'downwards'. It should be stressed that in general condition (4.3) is not sufficient for $\theta_t \in \arg_{\widehat{\theta}_t} \max W\left(\widehat{\theta}_t;\theta_t,\widehat{\theta}^{t-1}\right)$ to hold, though it certainly is necessary; the validity of the approach needs to be checked carefully in any given case.

Graphically, the potential pitfalls of the approach are illustrated using Figure 4.1. The vertical axis here denotes the value of $W\left(\cdot;\theta_t,\widehat{\theta}^{t-1}\right)$ for all given values of an agent's t-dated type report, which is mapped on the horizontal axis. To be sure that the incentive compatibility constraints are binding across all potential reports we would need to impose that this function is maximised at $W\left(\theta_t;\theta_t,\widehat{\theta}^{t-1}\right)$. Since this is an onerous requirement, as noted, our 'first-order' approach instead asserts simply that condition (4.3) must hold. We assume it is binding, and represent this by the horizontal line linking the value of $W\left(\cdot;\theta_t,\widehat{\theta}^{t-1}\right)$ at the relevant arguments in Figure 4.1. But knowing that $W\left(\cdot;\theta_t,\widehat{\theta}^{t-1}\right)$ does not change between θ_t' and θ_t is clearly not the same as knowing that $\theta_t \in \arg_{\widehat{\theta}_t} \max W\left(\widehat{\theta}_t;\theta_t,\widehat{\theta}^{t-1}\right)$. Whilst the rest of the value function certainly *may* be characterised by gradual and steady decay from the maximum, as in the case of the lower line AA in the figure, we equally cannot rule out the possibility of higher values being obtained elsewhere – as in the case of the line BB. In short, a stationary point need not imply a global maximum. The general implication (which extends to cases in which Θ is a continuum) is that the first-order approach admits a broader set of possible policies than the underlying incentive compatibility constraints, and unless one knows something about the properties of $W\left(\cdot;\theta_t,\widehat{\theta}^{t-1}\right)$ away from θ_t and θ_t' one can never be sure that a candidate policy satisfying condition (4.3) will *additionally* satisfy the full constraint set.

For this reason the following result is useful. The proof can be found in the appendix.

Proposition 1 ***Sufficiency of first-order approach:*** *Suppose the type set* Θ *contains only a finite number of elements and that under a given policy strategy the value function* $W\left(\widehat{\theta}_t;\theta_t,\widehat{\theta}^{t-1}\right)$ *satisfies increasing differences in* $\left(\widehat{\theta}_t,\theta_t\right)$*, so that the inequality*

$$W\left(\widehat{\theta}_t'';\theta_t'',\widehat{\theta}^{t-1}\right)-W\left(\widehat{\theta}_t';\theta_t'',\widehat{\theta}^{t-1}\right)>W\left(\widehat{\theta}_t'';\theta_t',\widehat{\theta}^{t-1}\right)-W\left(\widehat{\theta}_t';\theta_t',\widehat{\theta}^{t-1}\right)$$

holds for all $\left(\widehat{\theta}_t'',\widehat{\theta}_t',\theta_t'',\theta_t'\right)\in\Theta^4$ *such that* $\widehat{\theta}_t''>\widehat{\theta}_t'$ *and* $\theta_t''>\theta_t'$*. Then if condition (4.3) is known to hold with equality for all* $\theta_t\in\Theta\backslash\underline{\theta}$ *and all histories* $\widehat{\theta}^{t-1}\in\Theta^{t-1}$ *it must also be that* $W\left(\theta_t;\theta_t,\widehat{\theta}^{t-1}\right)>W\left(\theta_t'';\theta_t,\widehat{\theta}^{t-1}\right)$ *holds for all* $\theta_t''\in\Theta\backslash\left(\theta_t',\theta_t\right)$ *(where* $\theta_t'=\max\{\theta\in\Theta:\theta<\theta_t\}$ *and* $\underline{\theta}=\min\{\theta\in\Theta\}$*).*

This is a natural translation to our discrete-type setting of Theorem 5 in Kapička (2010).[13] Like that result, it is only an intermediate step in providing sufficiency conditions for the first-order approach, since the value function in any given setting will itself depend endogenously upon the chosen policy. But the result is nonetheless useful in supporting the arguments that follow. In words, it implies that a solution to the problem in which just condition (4.3) is imposed will also be a solution to the full problem (subject to the entire constraint set), provided the former solution exhibits the given increasing differences property. Moreover, combined with the single-crossing condition we have enough here to

[13] That theorem imposes that the derivative of W with respect to θ_t should be increasing in $\widehat{\theta}_t$.

assert something much stronger about the iid case, which we present in the following Corollary:

Corollary 1 *Suppose that agent-level productivities follow an iid process, and that the single-crossing condition (2.4) applies. Then provided a given policy strategy requires higher-type agents with a given history to produce higher output quantities than lower-type agents with the same history, and simultaneously provides them with higher consumption (in the period in which these productivities obtain), condition (4.3) holding with equality is sufficient for incentive compatability.*

Proof. When productivity shocks are iid, agents' future values (from $t+1$ on) for a given report $\widehat{\theta}^t$ are identical in expectation from the perspective of time t, irrespective of their true types. Hence increasing differences will follow provided we have, under the given policy:

$$u\left(\widehat{\theta}''_t;\theta''_t,\widehat{\theta}^{t-1}\right)-u\left(\widehat{\theta}'_t;\theta''_t,\widehat{\theta}^{t-1}\right)>u\left(\widehat{\theta}''_t;\theta'_t,\widehat{\theta}^{t-1}\right)-u\left(\widehat{\theta}'_t;\theta'_t,\widehat{\theta}^{t-1}\right)$$

where $u\left(\widehat{\theta}_t;\theta_t,\widehat{\theta}^{t-1}\right)$ is used to denote $u\left(c_t\left(\widehat{\theta}^t\left(\theta^\infty\right)\right),y_t\left(\widehat{\theta}^t\left(\theta^\infty\right)\right);\theta_t\right)$ for $\widehat{\theta}^t\left(\theta^\infty\right)=\left(\widehat{\theta}^{t-1},\widehat{\theta}_t\right)$, for all $\left(\widehat{\theta}''_t,\widehat{\theta}'_t,\theta''_t,\theta'_t\right)\in\Theta^4$ such that $\widehat{\theta}''_t>\widehat{\theta}'_t$ and $\theta''_t>\theta'_t$. The result is then a direct implication of the single crossing condition, given the assumption that output and consumption are increasing in type. ■

Whilst this result clearly still depends on the optimum having the particular property that output and consumption are increasing in type (for agents with a common reporting history), this is a very straightforward condition to check in any particular calculated example, and it will indeed generally hold under the optimal policy from the set satisfying condition (4.3).

In what follows we refer to the problem of policy choice from among the set of direct revelation mechanisms satisfying condition (4.3) as the 'relaxed' problem, in contrast with the 'general' problem that imposes $\theta_t\in\arg_{\widehat{\theta}_t}\max W\left(\widehat{\theta}_t;\theta_t,\widehat{\theta}^{t-1}\right)$ directly for all $\left(\theta_t,\widehat{\theta}^{t-1}\right)\in\Theta^t$.[14] Our focus will be on the properties of the solution to this relaxed problem, under the assumption that the solution to

[14] We distinguish between 'relaxed' and 'general' constraint sets in analogous fashion.

it coincides with the solution to the general problem. If this *is* the case, then we know that any other candidate policy that satisfies the constraint set of the relaxed problem is inferior from the policymaker's perspective to the solution to the general problem. We exploit this fact in what follows, showing how to perturb allocations in such a way that the constraint set of the relaxed problem must remain satisfied – and hence allocations must be inferior to the general problem's optimum.

5 Applying perturbation analysis

5.1 A diagramatic primer

This section introduces the main focal point of our analysis: how one can apply local perturbations to optimal consumption and output allocations in order to obtain a set of conditions that the optimal tax system must satisfy. The underlying innovation here is methodological, and the presentation builds up our new approach step by step. We additionally state important economic results relating to the character of optimal distortions – notably optimal savings wedges and optimal implicit income tax rates – whenever this is made possible by the general analysis. Since our focus is on direct revelation mechanisms rather than specific decentralisation schemes, these economic results must remain at a relatively high level of generality: they relate more to the *direction* in which any optimal tax wedges will distort allocations (relative to autarky) rather than the size and nature of specific taxes. Obviously it is hoped that further work will use the insight provided here to design effective decentralisation schemes.

The analysis that we devise is easiest to understand with the aid of an indifference curve map, linking output on the horizontal axis and consumption on the vertical. To make the relevant ideas concrete, and allow us to illustrate some important intuition for the dynamic tax problem, Figure 5.1 shows such a mapping.

The two within-period indifference curves are drawn for arbitrary distinct types θ_1 and θ_2, with

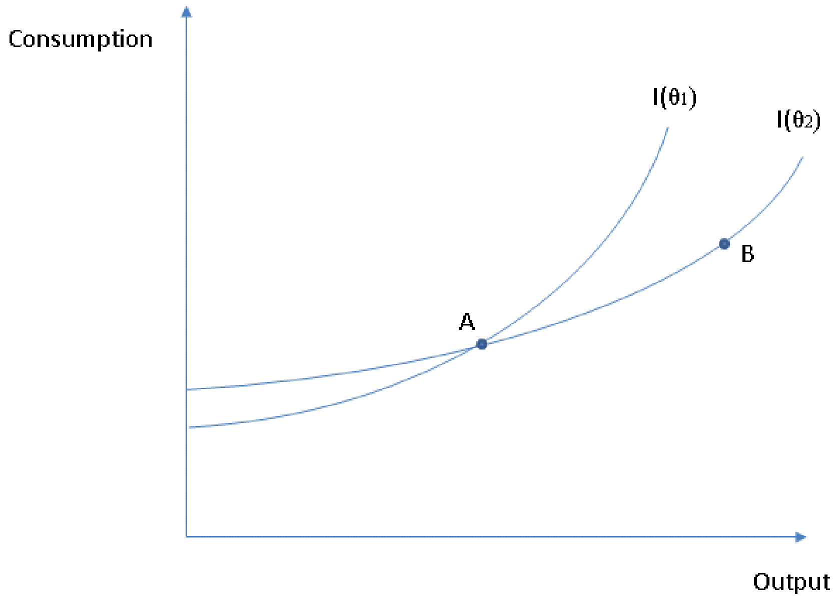

Figure 5.1: Graphical introduction to the policy trade-off

$\theta_2 > \theta_1$. The diagram can be used to show intuitively why positive effective marginal income taxes are desirable at a constrained optimum. Recall that the first-best allocation involved consumption-output bundles for each agent such that $u_c = -u_y$. Diagramatically this would correspond to a situation in which each agent's bundle is such that their indifference curve has a slope of 1, as if there is no taxation of income at the margin. We may suppose for illustrative purposes that this is true of points A (for the agent of type θ_1) and B (for type θ_2) in the diagram. We also know that at a first-best allocation the marginal utility of consumption would be equalised across agents (representable in the the case of separable utility by equalised consumption across all agents), and that there would be no history dependence in allocations.

When incentive compatibility constraints must additionally be satisfied these conditions can no longer be satisfied simultaneously. Figure 5.1 shows a situation in which the policymaker has chosen to violate just one of them: the equality of marginal utility across agents. Type θ_2 consumes at B, and is entirely compensated within the present period for choosing not to mimic type θ_1 (which would involve consumption at A). Yet this does not correspond to a second-best allocation. As the analysis of static optimal income tax problems has shown, the policymaker can improve matters by violating productive optimality for the *lower*-type agent. If type θ_1 is asked to produce at a point somewhere to the south-west of A along the curve $I(\theta_1)$, the 'information rent' that the higher-type agent can extract will be reduced. That is, the utility level type θ_2 could obtain by mimicking type θ_1 would fall, reducing the need for (wasteful) compensation – and thus freeing up resources to be redistributed to lower-type agents. Thus, perhaps counterintuitively, an equilibrium in which θ_1 is dissuaded from producing at the margin, via positive marginal income taxes, may be better for that agent than one in which there are zero marginal taxes.

Identical 'second best' logic may be applied to assert the desirability of spreading incentives through time. Rather than ensuring that the higher-type agent is just indifferent within a period between truthful reporting and mimicking, it will be preferable for that agent's within-period utility to be reduced – generating a strictly positive marginal benefit when the associated resources are redistributed – and for their discounted future utility instead to be raised in expectation by an offsetting amount. The latter distortion will come at zero initial marginal cost when one starts from a situation in which there is no history dependence – and so the theory of the second best applies: it will be better to introduce an extra dynamic distortion to mitigate the size of others.

5.2 Developing a perturbation approach

Our purpose is to make formal the intuition highlighted in the preceeding discussion. The presumption throughout is that the policymaker is able to solve the 'relaxed' problem, in which equation (4.3) replaces the complete constraint set, and that the solution to this problem coincides with the solution to the general problem in which the full constraint set is imposed – as well as implying strictly positive consumption and output allocations for all agents at all finite horizons.[15] Conditional upon a particular reporting history prior to the current period t, $\widehat{\theta}^{t-1}$, an agent's time-t report-contingent consumption and output allocations under the optimal scheme can be described by an $N \times 2$ matrix $X_t^*\left(\widehat{\theta}^{t-1}\right)$, with each row in this matrix corresponding to a particular $\widehat{\theta}_t \in \Theta$,[16] and the columns listing, in turn, consumption and output levels for the given reported productivity draw. Our aim is to show how these allocations can be perturbed by the addition of one or more of a particular set of $N \times 2$ matrices of continuously differentiable parametric functions, which in the generic case we denote by $\Delta(\delta)$ (with $\Delta : \mathbb{R} \rightarrow \mathbb{R}^{2N}$) for some relevant parameter δ (perhaps the consumption or utility increment implied by the given perturbation for an agent of the highest type). These functions are always normalised such that $\Delta(0) = 0$. In certain cases we will additionally allow changes to be spread through time, with the consumption and output of agents with a common reporting history $\widehat{\theta}^{t-1}$ changed at $t-1$ (as well as at t), according to an analogous function $\Delta_{-1}(\delta)$ (with $\Delta_{-1} : \mathbb{R} \rightarrow \mathbb{R}^2$). We wish to construct these Δ and Δ_{-1} functions so that they satisfy the following three properties:

1. Incentive compatibility constraints that bind under the relaxed problem for each time period up to the tth when allocations for agents with reporting history $\widehat{\theta}^{t-1}$ are $\left(c_{t-1}^*\left(\widehat{\theta}^{t-1}\right), y_{t-1}^*\left(\widehat{\theta}^{t-1}\right)\right)$at $t-1$ and $X_t^*\left(\widehat{\theta}^{t-1}\right)$ at t will continue to bind when allocations are $\left(c_{t-1}^*\left(\widehat{\theta}^{t-1}\right), y_{t-1}^*\left(\widehat{\theta}^{t-1}\right)\right) + \Delta_{-1}(\delta)$ at $t-1$ and $X_t^*\left(\widehat{\theta}^{t-1}\right) + \Delta(\delta)$ at t. Hence the perturbed allocations are candidate solutions to the relaxed problem.

2. $\Delta(\delta)$ and $\Delta_{-1}(\delta)$ should be both continuous and continuously differentiable in an open neighbourhood of the point $\delta = 0$.

3. Expected lifetime utility averaged across all agents should remain constant from the perspective of the initial time period for all values of δ in the neighbourhood of $\delta = 0$.

[15] Interiority of this kind is certainly essential to the reasoning we develop. It is what we intend in referring to an 'interior' solution to the relaxed problem in the analysis that follows.

[16] We assume that these are ordered in ascending values for $\widehat{\theta}_t$, with the lowest (reported) type's allocation in the first row of X^* and the highest type's in the Nth row.

Since we work under the assumption that incentive compatibility constraints bind only 'downwards' in the relaxed problem, the first property is equivalent to requiring that any additional incentive that an agent of type θ_t^n may have to mimic an agent of type θ_t^{n-1} (through changes in the allocation that the latter agent receives) is offset by an *equal* increase in the utility that the agent of type θ_t^n receives from truthful reporting.[17] Symmetrically, we impose that a reduction in the incentives to mimic should be matched by an equal reduction in the utility from truthful reporting – preserving continuity in the construction at $\delta = 0$. This ensures that if the original allocation satisfied the constraint set of the relaxed problem then the perturbed allocation must likewise. Hence if the original allocation was a *solution* to the relaxed problem, the perturbed allocation cannot deliver the same value to the policymaker at lower cost.

The second condition is required for the perturbations to be applied symmetrically. It is very similar to the requirement in consumer choice theory that optimal consumption should be at an interior point in an agent's budget set if we are to assert that the price ratio will be *equal* to that agent's marginal rate of substitution between two goods (and that a unique marginal rate of substitution should exist at the optimal point) – otherwise it may not be possible for the consumer to exploit any wedge that exists between the two. This requirement provides a substantial obstacle relative to the first: if we know that incentive compatibility constraints bind downwards then we know it always going to be possible to *increase* the utility of the highest type alone, or of the top n types in sufficiently skewed proportions, so that incentive compatibility constraints will remain satisfied. This could be done simply by the provision of extra consumption to higher-type agents. But perturbations of this form will only ever give us *in*equality restrictions – to the effect that the net marginal cost of changing outcomes in such a manner must be weakly positive. Unless a symmetric *downward* shift in the utility of high types is possible, with a converse impact on the net cost of utility provision, this cannot be converted into a first-order condition that is stated with *equality*.

[17] We use superscripts here to index the agents' types within the set Θ, with θ_t^n increasing in $n \in \{1, ..., N\}$

As the third condition states, we assume that allocations are changed in just such a way that the average value across agents of expected lifetime utility remains constant from the perspective of the very first time period. Since we have assumed a policymaker who is utilitarian, assessing outcomes from the perspective of the initial time period, this implies that in all cases the policymaker will experience no direct loss or gain from the perturbation.

A necessary condition for the original allocations $\left(c_{t-1}^{*}, y_{t-1}^{*}\right)$ and X_t^* to have been optimal is, then, that the marginal effect on the *resource cost* of utility provision associated with any admissible perturbation should be zero. Otherwise it would be possible to change allocations in one direction or another and raise a resource surplus, without changing the value of the policymaker's objective – contradicting optimality.

5.3 Deriving admissible perturbations: changes at the top

There is a very simple example of a perturbation that satisfies all three of the above requirements: a movement along the within-period indifference curve of the 'top' agent for any given reporting history. Since the famous work by Mirrlees (1971) it has been well understood that the maxim 'no distortion at the top' applies in a static optimal income tax setup – in the sense that $u_c = -u_y$ for any agent whose productivity parameter takes the highest possible value in the feasible set.[18] This derives from the fact that no other agent envies the allocation of the highest type in equilibrium – and thus there are no benefits in moving *away* from a situation in which $u_c = -u_y$.[19] The logic generalises to the intertemporal model, as the following makes clear.

Proposition 2 ***No distortion at the top:*** *Suppose the solution to the relaxed problem also solves the general problem, and is interior. Then in all time periods* $t \geq 1$ *and (if* $t > 1$*) for all past reporting*

[18] When this set has unbounded support the result need no longer hold, as the influential work by Saez (2001) has emphasised.

[19] For all other agents, reducing consumption and output together along a given indifference curve, to a point where $u_c > -u_y$, will reduce the utility 'rent' that must be provided to higher types to deter mimicking – a consideration that justifies deviating from the usual productive optimality condition in their case.

histories θ^{t-1}, the allocation (c_t^*, y_t^*) for the agent who reports θ_t' such that $\theta_t' = \max\{\theta \in \Theta\}$ satisfies $u_c\left(c_t^*, y_t^*; \theta_t'\right) = -u_y\left(c_t^*, y_t^*; \theta_t'\right)$.

Proof. Consider a perturbation to the allocation $X_t^*\left(\theta^{t-1}\right)$ given by the $N \times 2$ matrix of functions $\Delta : \mathbb{R} \to \mathbb{R}^{2N}$ such that the nth row of $\Delta(\delta)$ equals $(0, 0)$ for all $n \in \{1, ..., N-1\}$ and the Nth row equals $(\delta, f(\delta))$, with the function $f : \mathbb{R} \to \mathbb{R}$ defined implicitly by:

$$u\left(c_t^* + \delta, y_t^* + f(\delta); \theta_t'\right) = u\left(c_t^*, y_t^*; \theta_t'\right) \tag{5.1}$$

By construction this change keeps constant the (expected) utility of all truth-telling agents in all time periods. It does affect the utility of agents who report θ_t' when not of type θ_t', but this is irrelevant to the relaxed problem, and by the initial supposition in the Proposition we know that any allocation that continues to satisfy the relaxed constraint set cannot improve upon the solution to the general problem. This then implies that the value of the policymaker's objective will remain unchanged as δ is varied away from $\delta = 0$. The impact of the perturbation on the resources available to the policymaker in period t (in a truth-telling equilibrium) will be $\pi_\Theta\left(\theta_t'|\theta^{t-1}\right)\pi_\Theta\left(\theta^{t-1}\right)[f(\delta) - \delta]$. If the original allocation is optimal then the *marginal* impact on resources as δ moves away from zero must be zero, or else it would be possible to raise a surplus. Hence we have:

$$\pi_\Theta\left(\theta_t'|\theta^{t-1}\right)\pi_\Theta\left(\theta^{t-1}\right)\left[f'(0) - 1\right] = 0 \tag{5.2}$$

Probabilities are non-zero, so this implies:

$$f'(0) = 1 \tag{5.3}$$

Since utility for a highest-type truth-teller is unchanged by the perturbation we can assert the total derivative:

$$u_c\left(c_t^*, y_t^*; \theta_t'\right) + u_y\left(c_t^*, y_t^*; \theta_t'\right) f'(0) = 0 \tag{5.4}$$

The result follows immediately. ■

Notice that we have not had to assume anything about the type process in stating this proposition, which applies for any process consistent with the validity of the first-order approach. Graphically, the idea is that if the optimum involves only downwards incentive compatibility constraints binding then it must always be possible to move the allocation of the top agent at time t (for a given history) along that agent's within-period indifference curve, without jeopardising the incentives for any agent to report truthfully. This movement is additionally reversible, and (under the assumption of utilitarianism) will preserve the value of the policymaker's objective. Hence if the original allocation is optimal it must not raise surplus resources: the marginal cost of incentivising a top agent to produce an extra unit of output must exactly equal that extra unit.

The result is an interesting one in its own right, since Kocherlakota (2011) has provided a computed example in which the optimal consumption-output distortion for 'top' agents appears to be non-zero, conditional upon a particular past report.[20] Specifically, he obtains a non-zero 'top' rate in the second period of a two-period (overlapping generations) model for agents whose type was not the highest in the first period. The reason for this derives from the particular productivity process that he assumes. In the first period of his model, young agents may be either type θ_H (high type) of θ_L (low type). In the second period, those who were low types in the previous period may now be either type $\theta_L\theta_L'$ or type $\theta_L\theta_H'$, and those who were high types may be either type $\theta_H\theta_L'$ or type $\theta_H\theta_H'$. This implies that the highest type that an initially low-type agent could possibly be in the second period, $\theta_L\theta_H'$, is *not* the highest type across all agents in the economy, which is instead $\theta_H\theta_H'$. This in turn means that there are conceivably agents who could mimic the second-period agent of type $\theta_L\theta_H'$ whilst having

[20] See Chapter 6 of Kocherlakota (2011).

a productivity level in excess of $\theta_L\theta'_H$, as well as implying that two agents who receive the 'same' (stochastic component to their) productivity draw in the second period, θ'_H, do not have the same within-period preference structure across consumption-output space.

By contrast, in the model used in this paper the highest within-period type that an agent could *possibly* be is independent of history, and any two agents who receive the same within-period productivity draw and have reported the same history will make identical choices. Kocherlakota's results are influenced by the fact that changes to the second-period allocations of agents of type $\theta_L\theta'_H$ affect the incentives for first-period truthful reporting for agents of initial type θ_H (a point noted by the author). If we were to map from his setting to ours, the appropriate specification of Θ would be a time-varying set: $\Theta = \{\theta_L, \theta_H\}$ in the first period, and $\Theta = \left\{\theta_L\theta'_L, \theta_L\theta'_H, \theta_H\theta'_L, \theta_H\theta'_H\right\}$ in the second. It is only agents of type $\theta_H\theta'_H$ that we are claiming in this paper should see zero marginal rates in the second period, since $\theta_H\theta'_H$ is, in the relevant sense to us, the maximal element in Θ in the second period. This result (together with zero marginal rates for those of type θ_H in the first period) is indeed reported by Kocherlakota.

5.4 Uniform utility perturbations

The most common application of perturbation analysis in the dynamic optimal tax literature to date has been in deriving the 'inverse Euler condition' in models with additive separability in utility between consumption and labour supply. Deriving this condition relies on perturbing the *consumption* utility of certain agents in two consecutive periods. In the second of these, the consumption utility of all agents with a common prior report history is changed by a *uniform* amount at the margin. Because preferences are separable, under this perturbation agents will receive the same change to their within-period utility from mimicking any other agent as they do from a truthful report: separability implies consumption utility is type-independent. With non-separability we cannot construct perturbations that change the utility of all agents by an equal amount, no matter what their type. But we can appeal to the first-order approach, and focus just on making sure that the constraints of the *relaxed* problem remain satisfied. Then a natural generalisation of the inverse Euler condition to the non-separable case can be achieved, again based on a perturbation to allocations in two consecutive periods – and again ensuring that the utility of all agents with a common prior report history is changed by a *uniform* amount in the latter of these periods. When developing the analysis still further, we will see that this uniform utility provision is a special case of more general utility distributions that can be applied at the margin.

5.4.1 Definition of α function

Before presenting the main proposition of this sub-section, we must define the function $\alpha(c, y; \theta)$, with $\alpha : \mathbb{R}_+^2 \times \Theta \rightarrow \mathbb{R}$, as follows:

$$\alpha(c, y; \theta) = \frac{u_c(c, y; \theta) - u_c(c, y; \theta')}{u_y(c, y; \theta') - u_y(c, y; \theta)} \tag{5.5}$$

provided $\theta \neq \max\left\{\widetilde{\theta} \in \Theta\right\}$, where $\theta' = \min\left\{\widetilde{\theta} \in \Theta : \widetilde{\theta} > \theta\right\}$. If $\theta = \max\left\{\widetilde{\theta} \in \Theta\right\}$ we simply define $\alpha(c, y; \theta) = 0$.[21]

This α function is very useful in understanding the perturbation constructions that follow. It gives the marginal increase in output (away from the level y) that must accompany a unit marginal increase in consumption (away from the level c) *if the combined marginal perturbation is to have an equal impact on utility for the agents of both types,* θ *and* θ' at the given allocation.[22] More specifically for our purposes, it shows how to provide utility at the margin along a dimension in consumption-output space that will ensure both truth-tellers (θ-types) and would-be mimickers (θ'-types) receive the *same* utility increment.

If consumption is additively separable in utility then $\alpha = 0$ always holds. This is just a re-statement of the known result, used in deriving the standard inverse Euler condition, that the marginal effect of consumption changes on utility is completely independent of type under separability. In the general,

[21] Recall again that our focus at present is on the case in which Θ contains a finite number of elements.

[22] The impact of such a perturbation on the utility of type θ will be $u_c(c, y; \theta) + \alpha(c, y; \theta) u_y(c, y; \theta)$, and will be $u_c(c, y; \theta') + \alpha(c, y; \theta') u_y(c, y; \theta')$ for type θ'. It is easy to confirm that the two are equal.

non-separable problem it is not possible to find composite perturbations that have the same marginal effect on utility for *all* types in this way. But if it is sufficient to study the relaxed problem then the effects of perturbations really only matter to the extent that they change utility levels for *two* particular agents in each case: those truthfully reporting the given type, and would-be mimickers whose type is one higher. Moreover, it is always possible to ensure common utility changes for these two agents alone, even in the event of non-separability – and it proves useful to do so.

When consumption and labour supply are Edgeworth complements, so that higher levels of the latter increase the marginal utility of the former and vice-versa, we will have $\alpha > 0$.[23] That is, higher production must accompany higher consumption if the marginal increase in utility is to be the same for both truth-tellers and mimickers. This is because under complementarity the (truth-telling) lower-type agents will receive a greater marginal benefit from a unit increase in consumption at any given allocation than the (mimicking) higher-type agents – because of the higher number of hours the lower types are working to produce the given output level. To offset this disparity, one must exploit the higher marginal *dis*utility of additional output for lower types, by requiring that greater production should accompany the increased consumption. Conversely, when consumption and labour supply are Edgeworth substitutes we must have $\alpha < 0$.

5.4.2 A generalised inverse Euler condition

We now have the machinery to provide a generalisation of the inverse Euler condition to the case of non-separable preferences. Quite aside from its theoretical implications, this is of interest in its own right. On a simple analytical level, it helps fill a widely-recognised gap in the existing theory. Golosov, Tsyvinski and Werning (2006) have written that "Little is known about the solution of the optimal problem when preferences are not separable [between consumption and leisure]," before making use

[23] Formally, we take consumption and labour supply to be Edgeworth complements if and only if $u_{cy} > 0$, and Edgeworth substitutes if and only if $u_{cy} < 0$. Since these cross-partials hold θ fixed, higher output is equivalent to higher labour supply. Note that equation (2.3) further implies $u_{c\theta} < 0$ for Edgeworth complements and $u_{c\theta} > 0$ for Edgeworth substitutes.

of numerical simulations to show that some results (notably that savings 'wedges' should be positive) need not carry across from the separable to the non-separable case. Similarly, Kocherlakota (2011) has noted that "It would definitely be desirable to be able to construct optimal tax systems in dynamic settings in which preferences are nonseparable between consumption and labor inputs." The following result, it is hoped, will allow this to be achieved much more easily. The proof is slightly involved, but we choose to keep it in the main text because the methods used are novel and will be applied repeatedly throughout much of the subsequent analysis.

Proposition 3 ***Generalised inverse Euler condition:*** *Suppose the solution to the relaxed problem also solves the general problem, and is interior. Then for all time periods $t \geq 1$ and for all reporting histories θ^t, the allocations $\left(c_t^*\left(\theta^t\right), y_t^*\left(\theta^t\right)\right)$ and $X_{t+1}^*\left(\theta^t\right)$ satisfy the following condition:*

$$\begin{aligned} & R_{t+1}\beta\frac{1-\alpha\left(c_t^*, y_t^*; \theta_t\right)}{u_c\left(c_t^*, y_t^*; \theta_t\right)+u_y\left(c_t^*, y_t^*; \theta_t\right)\alpha\left(c_t^*, y_t^*; \theta_t\right)} \\ & = \sum_{\theta_{t+1}\in\Theta}\pi_\Theta\left(\theta_{t+1}|\theta^t\right)\frac{1-\alpha\left(c_{t+1}^*, y_{t+1}^*; \theta_{t+1}\right)}{u_c\left(c_{t+1}^*, y_{t+1}^*; \theta_{t+1}\right)+u_y\left(c_{t+1}^*, y_{t+1}^*; \theta_{t+1}\right)\alpha\left(c_{t+1}^*, y_{t+1}^*; \theta_{t+1}\right)} \end{aligned} \tag{5.6}$$

where c_{t+1}^* and y_{t+1}^* are given by the relevant entries in $X_{t+1}^*\left(\theta^t\right)$.[24]

Proof. We consider now a perturbation to outcomes in both time t and time $t+1$. Specifically, we wish to choose $\Delta\left(\delta\right)$ and $\Delta_{-1}\left(\delta\right)$ functions so that the agent with a truthful reporting history of θ^t will experience a reduction in within-period utility at time t of $\beta\delta$ units, and an increase in within-period utility at time $t+1$ of δ units *for any realisation of the $t+1$ productivity parameter*. These changes will, together, keep constant the expected utility associated with a truthful reporting strategy from the perspective of any time period up to the tth. The difficulty lies in constructing the perturbations in a way that will preserve incentive compatibility. Again, we exploit the supposition that no allocation that satisfies the constraint set of the relaxed problem can improve upon the solution to the general

[24] We suppress dependence upon past type reports to keep the notation manageable.

problem. This implies we need only concern ourselves with continuing to satisfy the $N-1$ constraints at $t+1$ that prevent mimicking by types 'one higher' than any given $\theta_{t+1} \in \Theta$, and the similar t-dated constraint preventing mimicking of type θ_t by the immediately superior type.

Indexing the elements of Θ in ascending order $\{1, ..., N\}$, our strategy is to construct perturbations in both time periods that change the consumption and output levels of the agent reporting θ^n in just such a way that the impact on within-period utility will be identical whether that agent is of true type θ^n or θ^{n+1}. To this end, let $\Delta_{-1}(\delta)$ be given by:

$$\Delta_{-1}(\delta) = (\phi^c(\theta_t, -\beta\delta; c_t^*, y_t^*), \phi^y(\theta_t, -\beta\delta; c_t^*, y_t^*)) \tag{5.7}$$

where $\phi^c(\theta, k; c^*, y^*)$ and $\phi^y(\theta, k; c^*, y^*)$ are defined implicitly when $\theta \neq \max\left\{\theta'' \in \Theta\right\}$ by the pair of equalities:

$$u(c^* + \phi^c(\theta, k; c^*, y^*), y^* + \phi^y(\theta, k; c^*, y^*); \theta) = u(c^*, y^*; \theta) + k \tag{5.8}$$

$$u\left(c^* + \phi^c(\theta, k; c^*, y^*), y^* + \phi^y(\theta, k; c^*, y^*); \theta'\right) = u\left(c^*, y^*; \theta'\right) + k \tag{5.9}$$

for $\theta' = \min\left\{\theta'' \in \Theta : \theta'' > \theta\right\}$, and when $\theta = \max\left\{\theta'' \in \Theta\right\}$ by

$$u(c^* + \phi^c(\theta, k; c^*, y^*), y^*; \theta) = u(c^*, y^*; \theta) + k \tag{5.10}$$

$$\phi^y(\theta, k; c^*, y^*) = 0 \tag{5.11}$$

That is to say, $\phi^c(\theta, k; c^*, y^*)$ and $\phi^y(\theta, k; c^*, y^*)$ are the consumption and output increments required to increase the utility of both mimickers and truth-tellers by k units. These functions will be uniquely defined, by the single crossing property. Similarly, the nth row of $\Delta(\delta)$is given by:

$$\left[\phi^c\left(\theta_{t+1}^n, \delta; c_{t+1}^*, y_{t+1}^*\right), \phi^y\left(\theta_{t+1}^n, \delta; c_{t+1}^*, y_{t+1}^*\right)\right] \tag{5.12}$$

where we index by type in the natural way. By construction this perturbation must preserve incentive compatibility in the relaxed problem at $t+1$, since the within-period utility that any agent can gain from mimicking is being changed by exactly the same amount (δ) as the within-period utility from truth-telling (for the mimicking strategies that need concern us). It must also preserve incentive compatibility at t under the relaxed problem, since its aggregate impact on the present value of expected utility from the perspective of period t and earlier is equal to zero (a reduction by $\beta\delta$ units at t and an increase by δ units at $t+1$, discounted at rate β), both for agents of true type θ_t and for the potential mimickers whose type is one higher. (Note that this assertion does not require any iid assumption, since utility is increased uniformly at the margin across all types at $t+1$.) The overall impact of the perturbation on the present value (assessed at time t) of the resources used by the policymaker is given by the following expression:

$$\begin{aligned}
&\pi_\Theta\left(\theta^t\right)\left[\phi^c\left(\theta_t, -\beta\delta; c_t^*, y_t^*\right) - \phi^y\left(\theta_t, -\beta\delta; c_t^*, y_t^*\right)\right] \\
&+R_{t+1}^{-1}\pi_\Theta\left(\theta^t\right)\sum_{\theta_{t+1}\in\Theta}\pi_\Theta\left(\theta_{t+1}|\theta^t\right)\left[\phi^c\left(\theta_{t+1}^n, \delta; c_{t+1}^*, y_{t+1}^*\right)\right. \\
&\left.-\phi^y\left(\theta_{t+1}^n, \delta; c_{t+1}^*, y_{t+1}^*\right)\right]
\end{aligned}$$

We require for optimality that the derivative of this expression with respect to δ should equal zero when $\delta = 0$; otherwise the policymaker could use fewer resources in obtaining the same value for aggregate utility, and still satisfy the relaxed problem's constraint set. Taking the derivative gives the optimality condition:

$$\begin{aligned}
&\beta\left[\phi_2^c\left(\theta_t, 0; c_t^*, y_t^*\right) - \phi_2^y\left(\theta_t, 0; c_t^*, y_t^*\right)\right] \\
= \; & R_{t+1}^{-1}\sum_{\theta_{t+1}\in\Theta}\pi_\Theta\left(\theta_{t+1}|\theta^t\right)\left[\phi_2^c\left(\theta_{t+1}^n, 0; c_{t+1}^*, y_{t+1}^*\right)\right. \\
&\left.-\phi_2^y\left(\theta_{t+1}^n, \delta; c_{t+1}^*, y_{t+1}^*\right)\right]
\end{aligned} \tag{5.13}$$

where ϕ_2^c denotes the derivative of ϕ^c with respect to its second argument. By total differentiation of conditions (5.8) to (5.11) with respect to k it is easy to show:

$$\phi_2^c(\theta, 0; c^*, y^*) - \phi_2^y(\theta, 0; c^*, y^*) = \frac{1 - \alpha(c^*, y^*; \theta)}{u_c(c^*, y^*; \theta) + u_y(c^*, y^*; \theta)\, \alpha(c^*, y^*; \theta)} \tag{5.14}$$

The result follows. ■

The important innovation here is to provide a general expression for the marginal cost of incentive-compatible utility provision from the perspective of the policymaker, and to show the manner in which it is optimal to spread this cost through time. Changing consumption and output jointly at t for the agent with report history θ^t according to the vector $(1, \alpha(c_t^*, y_t^*; \theta_t))$ increases the within-period utility of that agent at the margin by $u_c(c_t^*, y_t^*; \theta_t) + u_y(c_t^*, y_t^*; \theta_t)\, \alpha(c_t^*, y_t^*; \theta_t)$ units. By construction, it would have the same impact on a mimicking agent with a common report history to $t-1$, but a type one higher at t. The t-dated cost of providing utility in this manner at the margin for each agent with the given report history is $1 - \alpha(c_t^*, y_t^*; \theta_t)$ (any extra output being a negative cost). Hence the term on the left-hand side of (5.6) is the marginal cost for every β units of t-dated utility provided, which is converted into $t+1$ resources at the prevailing real interest rate. The term on the right-hand side is, by similar reasoning, the marginal cost (assessed at $t+1$) of providing the agent with report history θ^t with a guaranteed utility increment of one unit across types at time $t+1$ (and hence a discounted β units guaranteed from the perspective of t). Again, these marginal costs are obtained under the assumption that increments to a given $t+1$ type's utility must provide identical increases to the utility of mimicking agents.

Why do these marginal perturbations preserve incentive compatibility (at least for the relaxed problem)? Consider period $t+1$ first: we know that for any given agent the important consideration is whether the benefits to mimicking a type one lower have changed relative to the benefits from truthful reporting. This cannot be the case, since agents receive a common marginal utility increment of one unit in that time period whether they opt to be truth-tellers or 'downwards' mimickers. If truthful reporting was (weakly) optimal initially it must, therefore, remain so. This is the importance of assuming that output changes in accordance with the α function alongside any changes to consumption.

Meanwhile at time t the agent whose current type is indeed θ_t would see exactly offsetting changes to the present value of truth-telling were current utility to be increased (reduced) by an amount β at the margin and future utlity reduced (increased) across all future types by a unit at the margin. With no changes to the allocations to agents with alternative t-dated reports, this agent would have no reason not to continue reporting truthfully. But again, by construction we have ensured that the *same* (β-unit) marginal change to t-dated utility is engineered for the relevant *mimicker* (of type one higher than θ_t). And since a unit of utility is gained for all types at $t+1$ at the margin, this mimicker will likewise witness no change in the benefits to mimicking type θ_t. Hence incentive compatibility is preserved for any perturbation that increases (reduces) utility by β units at t and reduces (increases) it by one unit at $t+1$ for agents with the given reporting history. Since this perturbation is having no impact on lifetime utility for any agent from the perspective of period t and earlier, it must also be having no impact on the policymaker's objective function – so a necessary condition for optimality is that it cannot generate a surplus in net present value. This is what condition (5.6) is expressing.

Note that, like the 'no distortion at the top' condition, this result applies for general type processes – so long as the first-order approach remains valid.

5.4.3 Implications for optimal savings wedges

Generalising the inverse Euler condition in this way allows us to give qualified analytical support to the numerical result of Golosov, Tsyvinski and Werning (2006) that the optimal savings wedge *could* be negative for some agents, at least in the weak sense that under some preference structures we cannot analytically rule *out* the inequality:

$$u_c\left(c_t^*, y_t^*; \theta_t\right) > \beta R_{t+1} \sum_{\theta_{t+1} \in \Theta} \pi_\Theta\left(\theta_{t+1} | \theta^t\right) u_c\left(c_{t+1}^*, y_{t+1}^*; \theta_{t+1}\right) \tag{5.15}$$

holding in certain time periods for certain realisations of θ^{∞}. This would suggest tax instruments are being used to hold consumption at t *below* the level that would obtain in the event that the consumer could save freely at the gross real interest rate R_{t+1}, given the distribution of consumption across states in $t+1$; this can be interpreted as a marginal subsidy to savings. To understand why (5.15) may apply, it is worth recalling exactly why the optimal savings wedge is *positive* under separability.

Taking the mathematical treatment first, by Jensen's inequality we know that:

$$\sum_{\theta_{t+1}\in\Theta} \pi_{\Theta}\left(\theta_{t+1}|\theta^{t}\right)\left[\frac{1}{u_c\left(c^*_{t+1}, y^*_{t+1}; \theta_{t+1}\right)}\right] \geq \left[\sum_{\theta_{t+1}\in\Theta} \pi_{\Theta}\left(\theta_{t+1}|\theta^{t}\right) u_c\left(c^*_{t+1}, y^*_{t+1}; \theta_{t+1}\right)\right]^{-1} \tag{5.16}$$

with a strict inequality holding so long as the marginal utility of consumption varies in θ_{t+1} (which will be true in all models of interest). From this a simple rearrangement of (5.6) in the case of separable preferences ($\alpha\left(c, y; \theta\right) = 0$) confirms that savings are indeed deterred:[25]

$$u_c\left(c^*_t, y^*_t; \theta_t\right) \leq \beta R_{t+1} \sum_{\theta_{t+1}\in\Theta} \pi_{\Theta}\left(\theta_{t+1}|\theta^{t}\right) u_c\left(c^*_{t+1}, y^*_{t+1}; \theta_{t+1}\right) \tag{5.17}$$

(again, with a strict inequality holding so long as the marginal utility of consumption varies in θ_{t+1}).

The *economic* reason why the usual Euler condition (with an equality in the above relationship) does not hold in this environment derives from the linked problems of missing markets and over-insurance. Because each agent's productivity draw in each period is unobservable – and hence reports of it unverifiable – the idiosyncratic risk associated with future draws cannot be insured against. Absent any market intervention, the only way for individuals to prevent themselves from experiencing very low consumption in the event that they are unlucky in the future is to engage in private saving – their concerns dominated by future states of the world in which they are unlucky. This means all individuals in the economy are 'saving for a rainy day' simultaneously, even though it is (almost surely) impossible for them all to experience low productivity levels simultaneously. *Ex-post* there will

[25] See, for instance, Golosov, Kocherlakota and Tsyvinski (2003) for a fuller treatment of the separable utility case.

be a sizeable measure of individuals who were *not* unlucky, and thus who have a large quantity of accumulated savings that they would not have chosen to hold, had they been able access complete insurance markets. This excess stock of savings reduces the marginal benefits to these individuals from working, since the consumption returns from doing so are not that valuable to them. Thus over time more productive agents are content to put in less and less effort – an outcome that is not constrained efficient.[26] The policymaker prefers to rein in savings at the margin, making it easier to provide future production incentives for higher types.

A more direct way to understand the same result is simply to consider why the consumption Euler equation is not a necessary optimality condition for the policymaker's problem. The Euler condition states that spreading resources through time, with equal consumption increments across states at $t+1$, cannot raise utility. But this is not a choice open to our policymaker – who instead must ensure that spreading *utility* through time, with equal utility increments across states at $t+1$ (provided in a manner consistent with equal gains for truth-tellers and 'downwards' mimickers) cannot raise surplus *resources*. Providing equal consumption increments across states at $t+1$ would generally provide greater marginal utility to those whose initial consumption was lower, raising the benefits to higher types from mimicking them. In the separable case the marginal cost of utility provision *in a manner that preserves incentive compatibility* is the inverse of the marginal utility of consumption: only when consumption is provided *differentially* across states at $t+1$ in proportions according to this marginal cost can incentive compatibility be preserved. In the more general case this marginal cost is the expression contained in equation (5.6), with utility changes effected through a combination of consumption and output perturbations.

Perhaps slightly disappointingly, it turns out that a simple re-statement of inequality (5.17) in the non-separable case is only possible in very particular circumstances. Fortunately there is a natural

[26] There are clear parallels here with the general intuition provided by Greenwald and Stiglitz (1986) for missing markets and/or informational asymmetries implying a scope for Pareto-improving policy interventions (relative to market outcomes).

generalisation of the 'deterred savings' inequality that will apply more widely; but first we present the arguments that relate to this standard consumption Euler condition.

Proposition 4 ***Deterred savings (1):*** *Suppose the solution to the relaxed problem also solves the general problem and is interior, and additionally that in all time periods* $s \geq 1$ *and for all reporting histories* θ^s *the allocations* $(c_s^*(\theta^s), y_s^*(\theta^s))$ *imply* $u_c(c_s^*, y_s^*; \theta_s) \geq u_y(c_s^*, y_s^*; \theta_s)$*. Then for all time periods* $t \geq 1$ *and for all reporting histories* θ^t*, the allocations* $\left(c_t^*\left(\theta^t\right), y_t^*\left(\theta^t\right)\right)$ *and* $X_{t+1}^*\left(\theta^t\right)$ *will satisfy inequality (5.17) if one of the following conditions holds:*

1. *Preferences are additively separable between consumption and labour supply.*

2. *Consumption and labour supply are Edgeworth substitutes, and* $\theta_t = \max\{\theta \in \Theta\}$*.*

Proof. The result follows from the arguments above in the event that condition 1 holds, and does not require the extra assumption that intratemporal wedges are weakly positive.

The definition of Edgeworth substitutes gives $u_{cy} < 0$ under condition 2. By equation (2.3) this implies $u_{c\theta} > 0$. We also know $u_{y\theta} > 0$, so in general we will have $\alpha(c, y; \theta) \leq 0$, with a strict inequality except when $\theta = \max\left\{\theta' \in \Theta\right\}$. This, together with the assumption that intratemporal wedges are weakly positive, gives:

$$\frac{1 - \alpha(\theta)}{u_c(\theta) + u_y(\theta)\alpha(\theta)} \geq \frac{1}{u_c(\theta)} \tag{5.18}$$

(where we now suppress dependence upon c and y in the relevant functions to ease notation). Hence:

$$\begin{aligned} & \sum_{\theta_{t+1} \in \Theta} \pi_\Theta\left(\theta_{t+1} | \theta^t\right) \frac{1 - \alpha(\theta_{t+1})}{u_c(\theta_{t+1}) + u_y(\theta_{t+1})\alpha(\theta_{t+1})} \\ \geq & \sum_{\theta_{t+1} \in \Theta} \pi_\Theta\left(\theta_{t+1} | \theta^t\right) \frac{1}{u_c(\theta_{t+1})} \\ \geq & \left[\sum_{\theta_{t+1} \in \Theta} \pi_\Theta\left(\theta_{t+1} | \theta^t\right) u_c(\theta_{t+1})\right]^{-1} \end{aligned} \tag{5.19}$$

where the last result uses Jensen's inequality, and will hold strictly provided the marginal utility of consumption varies in θ_{t+1}. If $\theta_t = \max\{\theta \in \Theta\}$ then Proposition 2 implies $u_c(c_t^*, y_t^*; \theta_t) = u_y(c_t^*, y_t^*; \theta_t)$, so we have:

$$\begin{aligned}
R_{t+1}\beta\frac{1}{u_c(\theta_t)} &= R_{t+1}\beta\frac{1-\alpha(\theta_t)}{u_c(\theta_t)+u_y(\theta_t)\alpha(\theta_t)} \\
&= \sum_{\theta_{t+1}\in\Theta}\pi_\Theta\left(\theta_{t+1}|\theta^t\right)\frac{1-\alpha(\theta_{t+1})}{u_c(\theta_{t+1})+u_y(\theta_{t+1})\alpha(\theta_{t+1})} \\
&\geq \left[\sum_{\theta_{t+1}\in\Theta}\pi_\Theta\left(\theta_{t+1}|\theta^t\right)u_c(\theta_{t+1})\right]^{-1}
\end{aligned} \tag{5.20}$$

The result then follows from trivial manipulation. ■

We show subsequently that the assumption $u_c(c_s^*, y_s^*; \theta_s) \geq u_y(c_s^*, y_s^*; \theta_s)$ is indeed satisfied at any optimum: it is an immediate corollary of Proposition 8 below.

Thus we have a result that when consumption and labour supply are substitutes there will always be a positive savings wedge imposed on the highest-type agent, in the sense implied by inequality (5.17). Beyond this, though, it is hard to say much of specific relevance to the *consumption* Euler condition. But this condition isn't the *only* way to characterise a dynamically optimal decision in an economy free from government intervention. For instance, optimality under autarky also requires that a consumer cannot produce an extra unit of output at time t, save it, produce R_{t+1} units fewer at $t+1$, and increase the net present value of his or her utility by doing so. This consideration implies an alternative intertemporal optimality condition expressed in terms of an individual's output:

$$u_y(\theta_t) = \beta R_{t+1}\sum_{\theta_{t+1}\in\Theta}\pi_\Theta\left(\theta_{t+1}|\theta^t\right)u_y(\theta_{t+1}) \tag{5.21}$$

More significantly for our purposes, *any combination* of a reduction in consumption and increase in output at t, coupled with any distribution (in each state of the world) of the saved proceeds at $t+1$

between extra consumption and reduced output is possible, and similarly must not increase utility at an optimum under autarky (for resource movements towards either period). In particular, in a world with no taxation the following optimality condition must hold:

$$\begin{aligned} & \frac{u_c(\theta_t) + u_y(\theta_t)\,\alpha(\theta_t)}{1 - \alpha(\theta_t)} \\ = \; & \beta R_{t+1} \sum_{\theta_{t+1} \in \Theta} \pi_\Theta\left(\theta_{t+1} | \theta^t\right) \frac{u_c(\theta_{t+1}) + u_y(\theta_{t+1})\,\alpha(\theta_{t+1})}{1 - \alpha(\theta_{t+1})} \end{aligned} \tag{5.22}$$

The numerator in the object $\frac{u_c(\theta_t)+u_y(\theta_t)\alpha(\theta_t)}{1-\alpha(\theta_t)}$ is the marginal effect on the agent's utility at the given allocation of a unit increase in consumption, coupled with an increase in output of $\alpha(\theta_t)$ units. The denominator is the net cost to the agent of this change, under the maintained autarky assumption that all of the $\alpha(\theta_t)$ units of extra output are retained by the agent; the entire fraction then gives the marginal effect on utility per unit cost. The condition is just stating that no set of joint combinations of consumption and output changes can be used to spread resources through time and raise a surplus for the agent. So an agent's intertemporal ('savings') decisions are implicitly being distorted whenever equation (5.22) does not hold, with saving being discouraged whenever the left-hand side is less than the right.

Any such dynamic distortion may well interact with concurrent distortions at the labour-consumption margin *within* a period, but the main point here is that there is nothing inherently correct about focusing exclusively on deviations from the traditional *consumption* Euler equation in assessing the extent of savings distortions. Any equation that states that the marginal rate of substitution between a given pair of composite output-consumption bundles in two consecutive periods must equal their relative price ratio (in this case R_{t+1}), as equation (5.22) does, is of equal validity to the consumption Euler equation in characterising optimal dynamic behaviour under autarky.

The useful feature of equation (5.22) is that we can say something far more general about deviations from *this* expression at the optimum than we can about deviations from an Euler equation stated in terms of consumption alone. Specifically, we have the following.

Proposition 5 ***Deterred savings (2):*** *Suppose the solution to the relaxed problem also solves the general problem and is interior. Then for all time periods $t \geq 1$ and for all reporting histories θ^t, if consumption and labour supply are either Edgeworth substitutes or additively separable in preferences then savings will be deterred at the optimum, in the sense that the allocations $\left(c_t^*\left(\theta^t\right), y_t^*\left(\theta^t\right)\right)$ and $X_{t+1}^*\left(\theta^t\right)$ will satisfy the following condition:*

$$\begin{aligned} & \frac{u_c(\theta_t) + u_y(\theta_t)\alpha(\theta_t)}{1 - \alpha(\theta_t)} \\ \leq \quad & \beta R_{t+1} \sum_{\theta_{t+1} \in \Theta} \pi_\Theta\left(\theta_{t+1} | \theta^t\right) \frac{u_c(\theta_{t+1}) + u_y(\theta_{t+1})\alpha(\theta_{t+1})}{1 - \alpha(\theta_{t+1})} \end{aligned} \tag{5.23}$$

with the inequality holding strictly so long as the object $\frac{u_c(\theta_{t+1}) + u_y(\theta_{t+1})\alpha(\theta_{t+1})}{1 - \alpha(\theta_{t+1})}$ varies for different draws of $\theta_{t+1} \in \Theta$.

Proof. If consumption and labour supply are substitutes then $\alpha(\theta_t) < 0$, so for the preferences we are focusing on we must always have:

$$\frac{u_c(\theta_t) + u_y(\theta_t)\alpha(\theta_t)}{1 - \alpha(\theta_t)} > 0 \tag{5.24}$$

(recalling that $u_y < 0$). Thus by Jensen's inequality we have the following:

$$\begin{aligned} & \sum_{\theta_{t+1} \in \Theta} \pi_\Theta\left(\theta_{t+1} | \theta^t\right) \left[\frac{u_c(\theta_{t+1}) + u_y(\theta_{t+1})\alpha(\theta_{t+1})}{1 - \alpha(\theta_{t+1})}\right]^{-1} \\ \geq \quad & \left[\sum_{\theta_{t+1} \in \Theta} \pi_\Theta\left(\theta_{t+1} | \theta^t\right) \frac{u_c(\theta_{t+1}) + u_y(\theta_{t+1})\alpha(\theta_{t+1})}{1 - \alpha(\theta_{t+1})}\right]^{-1} \end{aligned} \tag{5.25}$$

with a strict inequality provided $\frac{u_c(\theta_{t+1}) + u_y(\theta_{t+1})\alpha(\theta_{t+1})}{1 - \alpha(\theta_{t+1})}$ varies for different draws of θ_{t+1}. The left-hand side of (5.25) is also the right-hand side of equation (5.6); the inequality in the Proposition then follows straightforwardly from using (5.6) in (5.25). ■

Note that this result has been stated irrespective of the manner and extent to which income is being taxed *within* periods t and $t+1$: unlike the prior Proposition we do not need any assumption that tax wedges are weakly positive for savings to be deterred in the given sense.

More significantly, note that we are *not* able to state the result for the case of Edgeworth complements: in that case we cannot rule out the possibility that $\frac{u_c(\theta_{t+1})+u_y(\theta_{t+1})\alpha(\theta_{t+1})}{1-\alpha(\theta_{t+1})} < 0$ may hold at the optimum for some values of θ_{t+1}, preventing us from applying Jensen's inequality. As it happens, Proposition 6 below implies $\frac{u_c(\theta_{t+1})+u_y(\theta_{t+1})\alpha(\theta_{t+1})}{1-\alpha(\theta_{t+1})} > 0$ will also hold at the optimum under complements provided types additionally follow an iid process for all agents; but this relies on arguments that we have yet to establish.

In economic terms, the result suggests the problem of over-saving in the absence of perfect insurance markets carries over directly to the case of substitutes. But we cannot be confident that savings will be being deterred at the margin *if the marginal cost of incentive-compatible utility provision could turn negative.* Though that possibility may at first appear unlikely, we show in a computed example below that it can indeed obtain under Markov shock processes and complementarity. The problem in this event is that, starting from the equilibrium allocation, an undistorted decision by agents to 'save' at the margin (as we have defined it) involves increasing their $t+1$ utility across all states through a uniform change in the quantity of resources at their disposal in that period, with these resources distributed between consumption and output in proportions corresponding to the relevant $\alpha(\theta_{t+1})$ terms. But it is possible that utility may be increased on average across $t+1$ states in this manner *even when the uniform quantity of resources to be allocated is negative.* This could happen whenever the optimum allows some agents to increase their utility despite increasing their production at the margin by more than their consumption – a possibility if there are large enough equilibrium distortions at the labour supply-consumption margin. In this case extra 'savings' – in the sense of incremental *utility* deferral – in fact correspond to a lower stock of *resources* being deferred. It would not be surprising if in this case the standard intuition relating to over-insurance did not apply.

This argument does highlight clearly the importance of distinguishing between marginal and average distortions. All we have said is that at the optimum under complementarity individuals could conceivably defer utility through constant marginal resource changes across future states, and pay a negative cost for doing so. But plainly this would never be a feature of the equilibrium allocation under autarky, which must always involve $u_c + u_y = 0$ for all agents in all periods – making it impossible for utility to increase at the margin along a vector that sees output rise by more than consumption. Knowing whether the *total* quantity of resources saved at the utilitarian optimum is less than the *total* quantity that would be saved under an autarkic equilibrium is as important as knowing whether the optimum is characterised by additional *marginal* savings being discouraged – but it is only the latter on which we have been able to shed light.

Moving away from its implications for marginal tax wedges, it will also be interesting to consider what Proposition 3 implies for the 'immiseration'-type results that emerge in the special case that $R_t = \beta^{-1}$ for all t. In that case equation (5.6) is a martingale, to which martingale convergence results may be applicable if bounds can be placed upon it. Under separable preferences the relevant martingale is in the inverse of the marginal utility of consumption, which is bounded below at zero under usual Inada conditions. It is well known (see, for instance, Farhi and Werning (2007)) that this implies almost all agents will see their marginal utility of consumption converge to the lower bound along an optimal path – and thus that consumption tends to zero for almost all agents. This 'immiseration' was first demonstrated as a potential property of optimal allocations under asymmetric information in a moral hazard context by Thomas and Worrall (1990), and it will turn out to generalise fairly robustly to the non-separable case – with important qualifications. But unfortunately the proofs rely on other arguments that are still to be established, so we defer a full treatment of this important area until later in this Part.

5.5 Intermediate perturbations: the case of iid types

5.5.1 Heuristic overview

We have shown above how it is possible to choose two particular pairs of $\Delta(\delta)$ and $\Delta_{-1}(\delta)$ schedules, in each case satisfying the three requirements of local incentive compatibility preservation, continuity, and zero impact on the policymaker's objective. The first was obtained by arguing that movements in either direction along the within-period indifference curve of the highest-type agent are always incentive-compatible (under the relaxed problem) and feasible. These perturbations will have zero impact on the within-period utility of all agents in the period that the $\Delta(\delta)$ schedule is applied (and in all other periods). The second was obtained by arguing we could reduce (increase) the utility of an agent with a given history to time t by $\beta\delta$ units according to $\Delta_{-1}(\delta)$, and increase (reduce) utility by δ units for all realisations of the productivity parameter at $t+1$ according to the matrix $\Delta(\delta)$ – in both cases in a manner that is incentive-compatible under the relaxed problem, and feasible. In both cases the results were quite general, in that they were derived without making any specific assumptions about the distribution of agent types through time.

This sub-section shows that it is also possible to find an entire set of perturbations 'intermediate' between these two extremes, in the sense that these can raise the within-period utility of just the top n types by δ units in the period that the relevant $\Delta(\delta)$ schedule is applied, whilst holding constant the utility of all other types (for all $n \in \{1, ..., N-1\}$). For the arguments to be valid we must make the additional simplifying assumption that type processes are iid for all agents. This is clearly unrealistic, but fortunately the logic generalises the Markov case with the addition of certain extra complications.

The intuition that we exploit is the following. Suppose one were to perturb the within-period allocation at time $t+1$ of an agent with prior reporting history $\widehat{\theta}^t$ and whose current report is $\widehat{\theta}^n_{t+1}$,[27] in a manner that keeps the utility of an agent truthfully reporting this type constant. This can be achieved by a movement along this agent's within-period indifference curve in consumption-output

[27] Superscripts again denote ranking in Θ, with higher values for n corresponding to higher values of θ.

space. If the solution to the relaxed problem also solves the general problem then we can can focus on changes that preserve incentive compatibility under the relaxed problem, knowing that they cannot improve on the best outcome when the general constraint set is imposed. To remain within the relaxed problem's constraint set we must simultaneously change the utility at $t+1$ of all *higher-type* agents (who share the same prior reporting history, $\widehat{\theta}^t$) by an amount equal to the change in the utility that the $n+1$th agent can now obtain by mimicking the nth. These latter utility changes must, in turn, be delivered along a dimension in consumption-output space consistent with the same impact being felt by truth-tellers and (relevant) mimickers – in the case of separability, for instance, via consumption increments alone.

To preserve prior incentive compatibility there must also be a perturbation in period t to the utility of an agent with reporting history $\widehat{\theta}^t$, so as to keep the present value of reporting $\widehat{\theta}^t$ constant in that period (and earlier) – with this perturbation again being constructed to ensure equal utility increments for truth-tellers and mimickers. The iid assumption will ensure that both truth-tellers and mimickers at t experience the same *ex-ante* change to their expected within-period utility at $t+1$ from reporting $\widehat{\theta}^t$, even though the proposed perturbation has differential effects across types at $t+1$.[28]

Figure 5.2 illustrates diagramatically the type of perturbation we have in mind for the latter period, under the assumption that there are two types in Θ. Suppose we know that the optimal allocation gives the lower-type agent a bundle at point A and the higher-type agent a bundle at point B in the relevant period.[29] Then suppose that the allocation to the lower type is perturbed along the lower type's indifference curve through A, $I(\theta_1)$. This will change the utility that the higher-type agent can obtain by mimicking, corresponding graphically to a movement by the higher type onto a distinct indifference curve from $I'(\theta_2)$ at the lower type's (perturbed) allocation, as movement takes place along $I(\theta_1)$. To offset this, the higher type must simultaneously be given a utility change when reporting

[28] Under more general type distributions this will not be true, since the probability distribution across future states will differ between mimickers and truth-tellers at t.

[29] As drawn this implies the higher-type agent strictly envies the lower type *within* the given period, so we are assuming higher future utility is also provided by way of compensation to prevent mimicking.

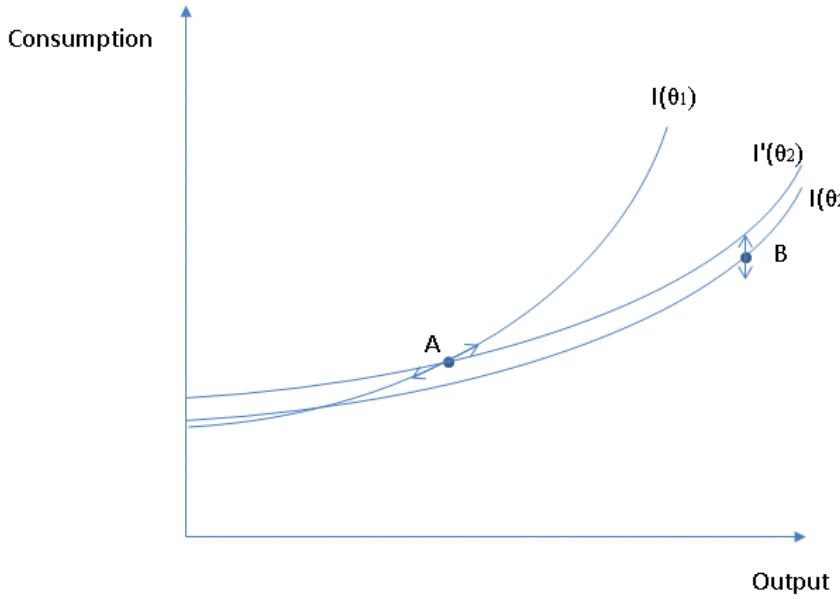

Figure 5.2: Preserving downwards incentive compatibility

truthfully, corresponding to a movement in the locus of point B by exactly the amount necessary for the higher-type agent to see the *same* change in utility at point A as at point B.

When more than two agents are present one must additionally provide still higher types with compensation for the given changes – but note that the within-period incentives of lower types will be unaffected. Any agent whose allocation is perturbed along his or her within-period indifference curve is, by construction, left no worse off from truthful reporting, and under the relaxed problem's constraint set this agent is not concerned by changes to allocations higher up in the type distribution. Still lower types are likewise unaffected. Provided we can additionally ensure incentives for truth-telling in prior periods are unaffected, this class of perturbation can be used to obtain additional optimality conditions.

5.5.2 Analytical treatment

We now present the arguments more formally. For each potential choice of n and any arbitrary report history $\widehat{\theta}^t$ for some $t \geq 1$, we wish to find a set of perturbations for period $t+1$, stacked in a function $\Delta(\delta)$, that will hold equilibrium consumption and output levels constant for all agents of type θ_{t+1}^{n-1} and lower, whilst changing them for type θ_{t+1}^{n} and higher.[30] This implies that the marginal impact on any given $\Delta(\delta)$ function of moving δ away from zero must depend on the choice of n, which in turn implies that if we are successful we will have obtained a new set of $N-1$ linearly independent dimensions along which the optimal allocation can be perturbed at the margin whilst preserving incentive compatibility locally – in turn providing $N-1$ new, distinct optimality conditions.

Before presenting the formal argument it is useful to define the function $\tau : \mathbb{R}_+^2 \times \Theta \to \mathbb{R}$ by $\tau(c, y, \theta) \equiv 1 + \frac{u_y(c,y,\theta)}{u_c(c,y,\theta)}$. This can be thought of as the implicit within-period marginal income tax rate faced by an agent of type θ receiving an allocation (c, y) – which is seen by noting:

$$u_c(c, y, \theta)(1 - \tau(c, y, \theta)) = -u_y(c, y, \theta) \tag{5.26}$$

Or, in words, the marginal utility of consumption multiplied by the real disposable income that an agent receives per unit of extra output that they produce is equal to the marginal disutility of production. By defining τ in this way we are not implying that a non-linear marginal income tax should necessarily form part of any ultimate decentralisation scheme, but it is useful to deploy a variable with a clear practical interpretation.

Then we have the following Proposition, the proof of which is in the appendix:

Proposition 6 ***Non-uniform utility increments:*** *Suppose the solution to the relaxed problem also solves the general problem and is interior, and that type draws are iid across agents and time. Then for all time periods $t \geq 1$, all reporting histories θ^t, and all $\theta_{t+1}^n \in \Theta : \theta_{t+1}^n \neq \max\{\theta' \in \Theta\}$ (so $n < N$),*

[30] Again, this applies only to the set of agents who have previously reported $\widehat{\theta}^t$.

the optimal allocations $\left(c_t^*\left(\theta^t\right), y_t^*\left(\theta^t\right)\right)$ *and* $X_{t+1}^*\left(\theta^t\right)$ *satisfy the following condition:*

$$
\begin{aligned}
& -\pi_\Theta\left(\theta_{t+1}^n\right) \frac{\tau\left(\theta_{t+1}^n\right)}{u_c\left(\widehat{\theta}_{t+1}^n;\theta_{t+1}^{n+1}\right)\left(1-\tau\left(\theta_{t+1}^n\right)\right)+u_y\left(\widehat{\theta}_{t+1}^n;\theta_{t+1}^{n+1}\right)} \\
& +\sum_{m=n+1}^{N} \pi_\Theta\left(\theta_{t+1}^m\right) \frac{1-\alpha\left(\theta_{t+1}^m\right)}{u_c\left(\theta_{t+1}^m\right)+u_y\left(\theta_{t+1}^m\right)\alpha\left(\theta_{t+1}^m\right)} \\
= \; & \beta R_{t+1}\pi_\Theta\left(\theta_{t+1}\geq\theta_{t+1}^{n+1}\right)\frac{1-\alpha\left(\theta_t\right)}{u_c\left(\theta_t\right)+u_y\left(\theta_t\right)\alpha\left(\theta_t\right)}
\end{aligned}
\tag{5.27}
$$

where $u\left(\widehat{\theta}_{t+1}^n;\theta_{t+1}^{n+1}\right)$ *is used to denote the within-period utility an agent of type* θ_{t+1}^{n+1} *can obtain by mimicking an agent of type* θ_{t+1}^n*, and its derivatives are defined correspondingly.*[31]

The term on the right-hand side of (5.27) is the marginal cost to the policymaker of raising utility at time t (for agents with the relevant report history) by an amount exactly equivalent to the discounted impact on expected utility of an increase in $t+1$ utility by a unit at the margin for all agents whose type is θ_{t+1}^{n+1} and higher. The term on the left-hand side is the marginal cost of raising $t+1$ utility in this fashion. The latter, in turn, has two components to it. The first is in fact a negative cost, so long as the effective marginal income tax rate is non-negative (and Proposition 8 below confirms that this is generally so). Any movement along the nth agent's indifference curve must raise resources so long as the slope of that indifference curve is less than one. The numerator $\tau\left(\theta_{t+1}^n\right)$ in the first fraction gives the net marginal increase in revenue to the policymaker (at $t+1$) for every unit increase in the consumption of agent θ_{t+1}^n, as we move along that agent's indifference curve.[32] The denominator (which is always positive) gives the marginal impact that this change has on the utility of mimickers whose type is one higher. Hence the entire term gives the additional quantity of resources raised from agents of type θ_{t+1}^n for every unit by which mimickers' utilities are increased. Notice that, normalised for scale, it must equal agent θ_{t+1}^n's marginal Hicksian (compensated) demand change for output minus that for

[31] We again suppress the dependence of τ, α, u_c and u_y on equilibrium allocations and report histories, in order to keep the notation manageable.

[32] To see this, recall that the slope of the relevant indifference curve is $-\frac{u_y\left(\theta_{t+1}^n\right)}{u_c\left(\theta_{t+1}^n\right)}=\left(1-\tau\left(\theta_{t+1}^n\right)\right)$

consumption as the effective price of consumption in terms of output is reduced at the equilibrium allocation.[33]

The second term on the left-hand side of (5.27) is more familiar from our earlier results: it gives the marginal cost of incentive-compatible utility provision to all agents above type θ_{t+1}^{n} – utility provision that is necessary for the complete perturbation to preserve incentive compatibility. The implications of the Proposition for this term are in fact best seen by applying (5.6), so as instead to obtain the following condition for all $n \in \{1, ..., N\}$ and $t \geq 1$:[34]

$$\begin{aligned} & -\frac{\pi_{\Theta}\left(\theta_{t+1}^{n}\right)}{\pi_{\Theta}\left(\theta_{t+1} \geq \theta_{t+1}^{n+1}\right)} \frac{\tau\left(\theta_{t+1}^{n}\right)}{u_c\left(\widehat{\theta}_{t+1}^{n}; \theta_{t+1}^{n+1}\right)\left(1-\tau\left(\theta_{t+1}^{n}\right)\right)+u_y\left(\widehat{\theta}_{t+1}^{n}; \theta_{t+1}^{n+1}\right)} \\ & +\sum_{m=n+1}^{N} \frac{\pi_{\Theta}\left(\theta_{t+1}^{m}\right)}{\pi_{\Theta}\left(\theta_{t+1} \geq \theta_{t+1}^{n+1}\right)} \frac{1-\alpha\left(\theta_{t+1}^{m}\right)}{u_c\left(\theta_{t+1}^{m}\right)+u_y\left(\theta_{t+1}^{m}\right) \alpha\left(\theta_{t+1}^{m}\right)} \\ = & \sum_{\theta_{t+1} \in \Theta} \pi_{\Theta}\left(\theta_{t+1}\right) \frac{1-\alpha\left(\theta_{t+1}\right)}{u_c\left(\theta_{t+1}\right)+u_y\left(\theta_{t+1}\right) \alpha\left(\theta_{t+1}\right)} \end{aligned} \tag{5.28}$$

Expressed in this way, we have a relationship between the expected marginal cost of utility provision across *all* types at $t+1$ (on the right-hand side), and the expected marginal cost conditional upon type being higher than θ_{t+1}^{n} (the second term on the left-hand side). Provided taxes are positive for the agent of type θ_{t+1}^{n}, the equation implies the expectation conditional upon being type θ_{t+1}^{n+1} or higher will be greater than the unconditional expectation. This is just a re-statement in the non-separable case of the fact that more resources must be used to increase by a given amount the welfare of those whose utility is already relatively high.

When utility is separable between consumption and labour supply, the previous condition collapses

[33] This follows directly from the fact that we are considering movements along the indifference curve.
[34] The case of $n = N$ reduces to the 'no distortion at the top' result.

to a much simpler object:

$$\begin{aligned} & \frac{\pi_{\Theta}\left(\theta_{t+1}^{n}\right)}{\pi_{\Theta}\left(\theta_{t+1} \geq \theta_{t+1}^{n+1}\right)} \frac{\tau\left(\theta_{t+1}^{n}\right)}{u_{y}\left(\widehat{\theta}_{t+1}^{n}; \theta_{t+1}^{n+1}\right) - u_{y}\left(\theta_{t+1}^{n}\right)} \\ & = E\left[\frac{1}{u_{c}\left(\theta_{t+1}\right)} | \theta_{t+1} \geq \theta_{t+1}^{n+1}\right] - E\left[\frac{1}{u_{c}\left(\theta_{t+1}\right)}\right] \end{aligned} \tag{5.29}$$

So at the optimum, the implicit tax rate will be higher the higher is the difference between the expected marginal cost of utility provision to agents *above* type θ_{t+1}^{n}, and its average across all agents. This makes sense: if one distortion from the first-best is high, in that high-type agents have relatively low marginal utilities of consumption, then other distortions are more likely to be beneficial. The effective tax rate will also be higher the greater is the productivity differential between truth-tellers of type θ_{t+1}^{n} and their mimickers, since high levels of this productivity gap imply distortions to the allocation of the lower type are more effective at the margin in 'pulling down' the utility from mimicking – and hence more desirable. Finally, the relative measure of agents of a higher type than θ_{t+1}^{n} will matter: the higher is this measure, the more desirable are higher effective taxes on θ_{t+1}^{n}, since the resources gained by reducing the utility of higher types are then more sizeable by comparison with the resources lost from distorting the production of the lower type.

All of these arguments are essentially familiar from the analysis of *static* optimal income tax models (see, in particular, Roberts (2000) and Saez (2001)) – though our focus on the marginal cost of utility provision is an innovation on that literature, and allows for a much simpler presentation of the the necessary optimality condition.[35] Indeed, though the arguments have not yet fully established it, we will subsequently show that condition (5.28) is necessary for intratemporal optimality regardless of whether earlier timer periods have existed, and this will be done in a manner that does not rely at all on *subsequent* time periods existing either – so that the condition must also apply at any optimum in a static income tax model (for which the first-order approach is valid). Though these models tend

[35] The simplification appears to arise from the fact our perturbations isolate Hicksian substitution effects by construction. I thank Kevin Roberts for highlighting this.

to assume a continuous type distribution, we show in a companion paper that the limit of (5.28) as types become arbitrarily close to one another is a necessary intratemporal optimality condition when the first-order approach is valid in that setting too (again under the.assumption of iid types) – see Brendon (2011). For this reason the arguments in the current paper could well prove equally useful in clarifying the analysis and implementation of static optimal income tax problems as of dynamic ones.

5.5.3 Characterising the set of perturbations

An interesting feature of condition (5.27) is that it nests our earlier two as extreme cases. Suppose we set $n = 0$, and extend the set Θ to include some arbitrary element, denoted θ^0, which is strictly less than all other elements of Θ and whose probability in any period is always zero under the π_Θ measure. Then $\pi_\Theta\left(\theta_{t+1}^0\right) = 0$ and $\pi_\Theta\left(\theta_{t+1} \geq \theta_{t+1}^{n+1}\right) = 1$ by definition, and the equation reduces to the generalised version of the inverse Euler equation, (5.6). If, on the other hand, we set $n = N$ then the last two terms drop out, and it reduces to a requirement that $\tau\left(\theta_{t+1}^N\right) = 0$ – the 'no distortion at the top' result from Proposition 2.

It proves useful for the analysis of the non-iid case to elaborate on this point further. Suppose $\phi^c(\theta, k; c, y)$ and $\phi^y(\theta, k; c, y)$ are defined for all $\theta \in \Theta$ as in the proof of Proposition 3, and $\varphi^c(\theta, k; c, y)$ and $\varphi^y(\theta, k; c, y)$ for all $\theta \neq \max\left\{\widetilde{\theta} \in \Theta\right\}$ as in the proof of Proposition 6.[36] If we define $\varphi^c(\theta, k; c, y)$ and $\varphi^y(\theta, k; c, y)$ for $\theta = \max\left\{\widetilde{\theta} \in \Theta\right\}$ by:

$$\varphi^c(\theta, k; c, y) = k \tag{5.30}$$

$$u\left(c + \varphi^c(\theta, k; c, y), y + \varphi^y(\theta, k; c, y); \theta\right) = u(c, y; \theta) \tag{5.31}$$

[36] Recall that $\phi^c(\theta, k; c, y)$ and $\phi^y(\theta, k; c, y)$ gave, respectively, the consumption and output changes necessary to increase the within-period utility of an agent of type θ by k units, starting from an allocation (c, y). For $\theta \neq \max\left\{\widetilde{\theta} \in \Theta\right\}$, $\varphi^c(\theta, k; c, y)$ and $\varphi^y(\theta, k; c, y)$ gave the corresponding consumption and output changes needed to increase the utility of a mimicking agent (one type higher than θ) by a unit, whilst holding constant the utility of θ types.

The derivatives of these functions with respect to their ith argument are denoted ϕ_i^c, ϕ_i^y etc.

and the derivatives with respect to k by φ_2^c and φ_2^y correspondingly, then the complete set of (non-marginal) $\Delta(\delta)$ perturbations that we consider at $t+1$ when obtaining Propositions 2, 3 and 6 is given by the $N+1$ matrices of the following form:

$$\begin{bmatrix} 0 & 0 \\ ... & ... \\ 0 & 0 \\ ... & ... \\ 0 & 0 \\ \varphi^c\left(\theta_{t+1}^N,\delta\right) & \varphi^y\left(\theta_{t+1}^N,\delta\right) \end{bmatrix}, \begin{bmatrix} 0 & 0 \\ ... & ... \\ 0 & 0 \\ ... & ... \\ \varphi^c\left(\theta_{t+1}^{N-1},\delta\right) & \varphi^y\left(\theta_{t+1}^{N-1},\delta\right) \\ \phi^c\left(\theta_{t+1}^N,\delta\right) & \phi^y\left(\theta_{t+1}^N,\delta\right) \end{bmatrix}$$

$$\begin{bmatrix} 0 & 0 \\ ... & ... \\ \varphi^c\left(\theta_{t+1}^n,\delta\right) & \varphi^y\left(\theta_{t+1}^n,\delta\right) \\ \phi^c\left(\theta_{t+1}^{n+1},\delta\right) & \phi^y\left(\theta_{t+1}^{n+1},\delta\right) \\ ... & ... \\ \phi^c\left(\theta_{t+1}^N,\delta\right) & \phi^y\left(\theta_{t+1}^N,\delta\right) \end{bmatrix}, \begin{bmatrix} \varphi^c\left(\theta_{t+1}^1,\delta\right) & \varphi^y\left(\theta_{t+1}^1,\delta\right) \\ \phi^c\left(\theta_{t+1}^2,\delta\right) & \phi^y\left(\theta_{t+1}^2,\delta\right) \\ ... & ... \\ \phi^c\left(\theta_{t+1}^{n+1},\delta\right) & \phi^y\left(\theta_{t+1}^{n+1},\delta\right) \\ ... & ... \\ \phi^c\left(\theta_{t+1}^N,\delta\right) & \phi^y\left(\theta_{t+1}^N,\delta\right) \end{bmatrix}$$

$$\begin{bmatrix} \phi^c\left(\theta_{t+1}^1,\delta\right) & \phi^y\left(\theta_{t+1}^1,\delta\right) \\ \phi^c\left(\theta_{t+1}^2,\delta\right) & \phi^y\left(\theta_{t+1}^2,\delta\right) \\ ... & ... \\ \phi^c\left(\theta_{t+1}^{n+1},\delta\right) & \phi^y\left(\theta_{t+1}^{n+1},\delta\right) \\ ... & ... \\ \phi^c\left(\theta_{t+1}^N,\delta\right) & \phi^y\left(\theta_{t+1}^N,\delta\right) \end{bmatrix}$$

to each of which corresponds a $\Delta_{-1}(\delta)$ perturbation given by

$$\left[\phi^c\left(\theta_t, -\beta\delta\pi_\Theta\left(\theta_{t+1} \geq \theta_{t+1}^{n+1}\right)\right), \phi^y\left(\theta_t, -\beta\delta\pi_\Theta\left(\theta_{t+1} \geq \theta_{t+1}^{n+1}\right)\right)\right]$$

where $n \in \{1, ..., N+1\}$ identifies the lowest agent type at $t+1$ whose utility is being increased (with $N+1$ corresponding to movements along the 'top' indifference curve only). In what follows we index the entire matrix of perturbations by this n, so that the first of the matrices above gives $\Delta^{N+1}(\delta)$, with corresponding $\Delta_{-1}^{N+1}(\delta)$:

$$\Delta_{-1}^{N+1}(\delta) = [0, 0] \tag{5.32}$$

the second gives $\Delta^N(\delta)$, with corresponding $\Delta_{-1}^N(\delta)$:

$$\begin{aligned} &\Delta_{-1}^N(\delta) \\ = \; &\left[\phi^c\left(\theta_t, -\beta\delta\pi_\Theta\left(\theta_{t+1} \geq \theta_{t+1}^N\right)\right), \phi^y\left(\theta_t, -\beta\delta\pi_\Theta\left(\theta_{t+1} \geq \theta_{t+1}^N\right)\right)\right] \end{aligned} \tag{5.33}$$

and so on. The last matrix, corresponding to uniform utility provision across agents, we denote $\Delta^1(\delta)$, with corresponding $\Delta_{-1}^1(\delta)$:

$$\Delta_{-1}^1(\delta) = \left[\phi^c(\theta_t, -\beta\delta), \phi^y(\theta_t, -\beta\delta)\right] \tag{5.34}$$

The *marginal* effects on allocations of moving δ away from zero will, for each of these perturbations, be given by equivalent matrices in which $\phi_2^c(\cdot, 0)$ and $\phi_2^y(\cdot, 0)$ replace $\phi^c(\cdot, \delta)$ and $\phi^y(\cdot, \delta)$ respectively, and $\varphi_2^c(\cdot, 0)$ and $\varphi_2^y(\cdot, 0)$ replace $\varphi^c(\cdot, \delta)$ and $\varphi^y(\cdot, \delta)$. It is then clear by inspection that the $N+1$ marginal changes to allocations at $t+1$, which we may denote $\Delta^{n\prime}(0)$ for $n \in \{1, ..., N+1\}$, are linearly independent from one another. Thus each of the associated first-order conditions will be providing distinct information about the character of the optimal allocation.

The sense in which the perturbations analysed in Proposition 6 are 'intermediate' between those of Propositions 2 and 3 should now be apparent. A compact way to see matters is in terms of an $(N+1) \times N$ matrix listing the set of possible marginal impacts on the within-period utility levels of agents of different types that is afforded by the marginal perturbations $\Delta^{n\prime}(0)$ for $n \in \{1, ..., N+1\}$. We label this matrix $\widehat{J}$:

$$\widehat{J} = \begin{bmatrix} 1 & 1 & \cdots & 1 & 1 \\ 0 & 1 & \cdots & 1 & 1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & 1 \\ 0 & 0 & \cdots & 0 & 1 \\ 0 & 0 & \cdots & 0 & 0 \end{bmatrix} \tag{5.35}$$

Element $\widehat{J}_{n,m}$ is then the marginal impact on the within-period utility of an agent of type θ_{t+1}^{m} caused by a marginal move in δ away from zero in accordance with the $\Delta^{n}(\delta)$ schedule. Note the fact that the non-zero entries are all equal to one arises from the way we have defined δ in each case. In what follows it is useful to work with the $N \times N$ square matrix obtained by deleting the last (zero) row from $\widehat{J}$. We label this square matrix J.

This matrix representation allows us to develop the analysis in a way that will prove very helpful when we move away from iid types, and we next present two technical Lemmas that are useful to this end. In stating them we denote by π_{Θ}^{vec} the N-dimensional column vector stacking the probabilities $\pi_{\Theta}(\theta_t^n)$ in order from $n = 1$ to $n = N$ (these are independent of history and time under the maintained iid assumption). Then we have the following, the proof of which is in the appendix:

Lemma 1 *Suppose that type draws are iid across agents and time. Then for any vector $\nu \in \mathbb{R}^N$ (whose nth element is denoted ν_n), all time periods $t \geq 1$ and any given reporting history θ^t, it is possible to perturb the set of optimal allocations $\left(c_t^*(\theta^t), y_t^*(\theta^t)\right)$ and $X_{t+1}^*(\theta^t)$ (assumed to be interior) in a manner that will preserve the incentive-compatibility constraints of the relaxed problem in all periods whilst raising the within-period utility of an agent of type θ_{t+1}^n by an amount $\nu_n\delta$ at $t+1$ and raising* *the within-period utility of the agent at t by an amount $-\beta\nu'\pi_{\Theta}^{vec}\delta$, for any scalar δ satisfying $|\delta| < \varepsilon$ for some $\varepsilon > 0$. Additionally, considering period 1 in isolation, for all vectors $\nu \in \mathbb{R}^N$ that satisfy $\nu'\pi_{\Theta}^{vec} = 0$ it is possible to perturb the allocations X_1^* in a manner that will preserve the incentive-compatibility constraints of the relaxed problem whilst raising the within-period utility of an agent of type θ_1^n by an amount $\nu_n\delta$ in period 1, again for any scalar δ satisfying $|\delta| < \varepsilon$ for some $\varepsilon > 0$.*

This result is useful for two reasons. First, because it implies a set of 'intratemporal' optimality conditions that must hold in the initial period, which it was not possible to obtain through a proof that appealed along the way to perturbations to allocations in a *prior* time period.[37] Second, because for each 'utility increment' vector ν there will always be an equivalent $N \times 1$ vector γ, satisfying:

$$\gamma = \left(J^{-1}\right)' \nu \tag{5.36}$$

Since J^{-1} is plainly invertible, the Lemma implies that we can pick composite sums of the marginal utility changes that are induced by each of the $\Delta^{n\prime}(0)$ marginal perturbations (for $n \in \{1, ..., N\}$), in proportions corresponding to any $\gamma \in \mathbb{R}^N$ (that is, γ_1 units of the utility change from marginal perturbation $\Delta^{1\prime}$, plus γ_2 units of $\Delta^{2\prime}$, and so on), and know that the overall vector of utility changes described in this way, given by $J'\gamma$, can be implemented through some incentive-compatible perturbation (both at the margin and for discrete utility increments).

Moreover, if one solves for the marginal effect on the policymaker's surplus at $t+1$ that is associated with the perturbation that changes within-period utility away from the optimum according to the vector ν (per unit increase in δ), we can also show that this is equal to the effect on the policymaker's surplus of an additive combination of the $\Delta^{n\prime}(0)$ marginal perturbations for $n = \{1, ..., N\}$, in proportions that likewise correspond to the entries in the associated γ. Together with movements along the Nth agent's indifference curve (which have no effect on the utility of any agent and thus cannot assist in the provision of utility according to any vector ν), the complete set of $\Delta^{n\prime}(0)$ marginal perturbations for

[37] This is the purpose of stating separately the final sentence in the Proposition.

$n \in \{1, ..., N+1\}$ will therefore *span* the entire set of $(N+1)$ dimensions along which the consumption and output of all agents can be jointly perturbed at the margin away from the optimal allocation X_{t+1}^* without undermining the within-period incentive compatibility requirements of the restricted problem at $t+1$. This is very useful for technical purposes, since it implies any complex marginal perturbation to X_{t+1}^* that is known to satisfy within-period incentive compatibility at $t+1$ can always be written as a composite of the $\Delta^{n\prime}(0)$ matrices – complementing the previous Lemma, which implied that any composite of the $\Delta^{n\prime}(0)$ matrices was within-period incentive-compatible.

We summarise the result in the following Lemma, the proof of which is again relegated to the appendix.

Lemma 2 *Consider perturbations of the form outlined in Lemma 1, generating utility changes at $t+1$ according to the vector $\delta\nu$. For any such $\nu \in \mathbb{R}^N$ the marginal resource cost of this perturbation to the policymaker at $t+1$ can always be expressed as the following additive combination of the $\Delta^{n\prime}(0)$ matrices for $n \in \{1, ..., N\}$:*

$$\pi_\Theta\left(\theta^t\right)\left(\pi_\Theta^{vec}\right)'\left[\sum_{n=1}^{N}\gamma_n\Delta^{n\prime}(0)\right]k$$

where k is here defined as the 2×1 vector $\begin{bmatrix} 1 \\ -1 \end{bmatrix}$ and γ_n is the nth element in the vector γ, which in turn satisfies:

$$\gamma = \left(J^{-1}\right)'\nu \tag{5.37}$$

Moreover, any perturbation to a constrained-optimal allocation that changes allocations at t and $t+1$ across agents with a common reporting history to t whilst remaining within the constraint set of the relaxed problem must have marginal effects on allocations at $t+1$ that are expressible as a linear combination of the $\Delta^{n\prime}(0)$ marginal effects alone.

The proof of this Lemma does not rely on the iid assumption: if a given perturbation to utilities admissible under Lemma 1 is additionally incentive-compatible under a more general Markov type process then it will also yield a surplus according to the expression above.

Together these results allow us to provide a more general statement of the requirements for 'intratemporal' optimality in the iid case – in the style of condition (5.28) above. This will enable us to divide up the set of $N+1$ optimality conditions that we have derived into N *intra*temporal conditions, and just one *inter*temporal condition. Because an intermediate step in our original derivation of condition (5.28) was the use of utility perturbations in the *preceeding* time period, at present we strictly can only state it for periods after the first. With the aid of Lemma 1 this can easily be overcome, as the next Proposition shows. First, let γ be any $N \times 1$ vector with the following property:

$$\gamma' J \pi_{\Theta}^{vec} = 0 \tag{5.38}$$

Then we have the following result.

Proposition 7 ***Intratemporal optimality (iid case):*** *Suppose the solution to the relaxed problem also solves the general problem and is interior, and that type draws are iid across agents and time. Then for all time periods $t \geq 1$ and all reporting histories θ^t, the matrices $\{\Delta^{n\prime}(0)\}_{n=1}^{N}$ associated with the optimal allocation matrix $X_{t+1}^{*}(\theta^t)$ satisfy the following condition:*

$$(\pi_{\Theta}^{vec})' \left[\sum_{n=1}^{N} \gamma_n \Delta^{n\prime}(0) \right] k = 0 \tag{5.39}$$

where k is again the 2×1 vector $\begin{bmatrix} 1 \\ -1 \end{bmatrix}$ and γ_n is the nth element of any vector γ that satisfies equation (5.38).

Proof. Appealing to Lemma 1, we know that the vector ν that solves $\nu = J'\gamma$ can be applied at the margin to augment the utilities of agents with a common prior type history in any period $t \geq 1$, through a change to output and consumption bundles that preserves incentive compatibility. By the definition of γ we will have $\nu' \pi_{\Theta}^{vec} = 0$, and thus incentive compatibility in any prior periods is assured without the need for any further perturbations. The result then follows directly from Lemma 2. ■

Notice that there will, in general, exist $N-1$ linearly independent γ vectors satisfying equation (5.38). Equation (5.28) (which we may now assert for all time periods, including the first) gives the $N-1$ possibilities for which $\gamma_1 = -1$, $\gamma_n = [\pi_{\Theta}(\theta_t \geq \theta_t^n)]^{-1}$ for some $n \in \{2, ..., N\}$ and any $t \geq 1$, and there are zero entries elsewhere. When we move to the non-iid case the set of admissible γ vectors is reduced in an important way, and this matrix representation proves invaluable in characterising composite movements that *are* still possible in that case.

Recall that there will additionally exist an Nth intratemporal condition – the 'no distortion at the top' result – to which any set of $N-1$ linearly independent composite perturbations that satisfy the requirements of Proposition 7 should be added. For all time periods after the first, the generalised inverse Euler condition, (5.22), provides a further cross-restriction, linking outcomes for agents with a given prior history to their common allocation in the previous period. There are additionally $N-1$ binding incentive compatibility constraints across any N agents who share a common prior history. Finally, there is a single intertemporal budget constraint that the policymaker must satisy (which may be thought of loosely as 'substituting' for the dynamic Euler condition in the very first time period) – ensuring that we always have precisely $2N$ equations to tie down the $2N$ variables that are to be determined across any set of agents with a common prior history at any given point in time. In this sense we have provided a complete analytical description of the solution. This is likely to be of great use practically, since it obviates the need to apply dynamic programming techniques in arriving at a numerical solution to any given example. In a finite-horizon model, all that is needed is to solve a known set of simultaneous equations, although in general the number of equations will grow exponentially in the number of time periods.[38]

[38] For T periods there will generally be $\sum_{t=1}^{T} 2N^t$ variables to determine.

In an infinite-horizon model one will still need a way to approximate agents' value functions conditional on any shock history (since these feature in the binding incentive compatibility constraint). But note that in the iid case history dependence can be summarised by a single variable – the marginal cost of uniform incentive-compatible utility provision. This follows from the fact that we can solve for outcomes from period t onwards for agents with a given type history θ^{t-1} based on a set of optimality conditions and constraints that all depend only on outcomes from t onwards, plus the generalised inverse Euler condition applied between $t-1$ and t. The marginal cost of utility provision must be sufficient to summarise the past completely.

5.5.4 Optimal effective income tax rates

The results of the previous subsection allow us to demonstrate a further quite general result with important economic implications.

Proposition 8 ***Non-negative income taxes:*** *Suppose the solution to the relaxed problem also solves the general problem and is interior. Then for all time periods $t \geq 1$, all reporting histories θ^{t-1} and all $\theta_t^n \in \Theta$ the implicit marginal tax rate $\tau(\theta_t^n)$ satisfies $\tau(\theta_t^n) \geq 0$.*

Proof. Consider the perturbation given by applying just the nth row of the schedule $\Delta^{n+1}(\delta)$ at time t – that is, a movement along the within-period indifference curve of the nth agent. For negative values of δ (only) this will keep us within the constraint set of the relaxed problem, since the net impact on the utility obtainable from reporting θ^t at t is zero for truth-tellers and strictly negative for 'downwards mimickers', and expected utility in prior periods is left completely unaffected by the fact that all agents at t are indifferent to this perturbation. Hence the marginal cost as δ is moved marginally *below* zero must be weakly positive, given that the optimal solution in the relaxed constraint set solves the general problem. From our earlier results, this implies:

$$-\left[\varphi_2^c\left(\theta_{t+1}^n, 0\right) - \varphi_2^y\left(\theta_{t+1}^n, 0\right)\right] \geq 0 \tag{5.40}$$

The proof of Proposition 6 shows:

$$\varphi_2^c\left(\theta_t^n, 0\right) - \varphi_2^y\left(\theta_t^n, 0\right) = -\frac{\tau\left(\theta_t^n\right)}{u_c\left(\widehat{\theta}_t^n; \theta_t^{n+1}\right)\left(1-\tau\left(\theta_t^n\right)\right) + u_y\left(\widehat{\theta}_t^n; \theta_t^{n+1}\right)} \tag{5.41}$$

where $u\left(\widehat{\theta}_t^n; \theta_t^{n+1}\right)$ (and associated partial derivatives) denotes the utility function of an agent whose type is θ_{t+1}^{n+1} mimicking one of type θ_{t+1}^n. Hence:

$$\frac{\tau\left(\theta_{t+1}^n\right)}{u_c\left(\widehat{\theta}_{t+1}^n; \theta_{t+1}^{n+1}\right)\left(1-\tau\left(\theta_{t+1}^n\right)\right) + u_y\left(\widehat{\theta}_{t+1}^n; \theta_{t+1}^{n+1}\right)} \geq 0 \tag{5.42}$$

We have:

$$\begin{aligned}\left(1-\tau\left(\theta_{t+1}^n\right)\right) &= -\frac{u_y\left(\theta_{t+1}^n\right)}{u_c\left(\theta_{t+1}^n\right)} \\ &> -\frac{u_y\left(\widehat{\theta}_{t+1}^n; \theta_{t+1}^{n+1}\right)}{u_c\left(\widehat{\theta}_{t+1}^n; \theta_{t+1}^{n+1}\right)}\end{aligned} \tag{5.43}$$

where the last inequality is an application of the single-crossing condition. Hence the denominator in condition (5.42) will be strictly positive, and the result follows. ■

So unlike the savings distortion the direction of the intratemporal distortion on production is completely unambiguous: the optimal effective marginal income tax rate is never negative. Note that we have not had to make any iid assumption in stating this result. This follows from the fact the the chosen perturbation does not change equilibrium utility levels across agents, so differences in previous periods between the probability distributions of truth-tellers and mimickers are irrelevant to continued incentive compatibility.

In a sense the result itself should not be surprising. We have already seen that the first-best involves effective marginal tax rates of zero on current income, and there are benefits to moving away from this situation under imperfect information only to the extent that doing so reduces the information rent that higher types are able to extract as compensation for not mimicking. This was the message of Figure 5.1 above. Since a 'downwards' movement along the within-period indifference curve of lower types reduces the utility of higher-type mimickers, it is always better to move to a point where this indfference curve has a slope ($\frac{dc}{dy}$) that is less than one.

6 General perturbations with Markov types

Whilst the iid model is instructive, it is plainly unrealistic as a description of the way individuals' productivities evolve in practice. To attain some greater realism we need to generalise to allow for persistence in types. The simplest way to do this is to assume the productivity measure π_Θ incorporates a Markov structure (so $\pi_\Theta\left(\theta_{t+1}|\theta^t\right) = \pi_\Theta\left(\theta_{t+1}|\theta_t\right)$). Recall from the earlier discussion that our confidence in the first-order approach cannot be so sure in this case: we had to assume increasing differences in the value function at the relaxed problem's optimum for sufficiency, which was not a condition directly related to the 'fundamentals' of the model. We proceed all the same, and leave a more satisfactory resolution of the sufficiency question for subsequent work.

When types follow a general Markov process we are faced with an extra dimension of complication. For agents with a given reporting history $\widehat{\theta}^{t-1}$ we may be able to define a perturbation to allocations at t that has zero impact on the expected utility at $t-1$ of a relevant *truth-telling* agent, but the probability distribution under which this expectation is calculated is now particular to that agent. An agent who is, at the optimum, on the cusp of *falsely* reporting $\widehat{\theta}^{t-1}$ will take expectations of the future returns from a mimicking strategy under a different probability distribution to the truth-teller – and thus may well experience a change in the *ex-ante* expected utility from mimicking subsequent to the perturbation eveb though the truth-teller does not. This would undermine local incentive compatibility at time $t-1$, for movements in one direction or the other.

In general our aim is, once again, to find a set of distinct functions $\Delta : \mathbb{R} \rightarrow \mathbb{R}^{2N}$ and $\Delta_{-1} : \mathbb{R} \rightarrow \mathbb{R}^2$ that can be used to perturb the consumption and output allocations across all agents with a given reporting history $\widehat{\theta}^{t-1}$, at t and $t-1$ respectively, subject to these functions satisfying the three conditions set out at the start of Section 5.2: the preservation of incentive compatibility, continuous differentiability in δ in the region of $\delta = 0$, and no net impact on the policymaker's initial-period objective. It is the first of these conditions – incentive compatibility – that we will no longer necessarily satisfy through applications of the Δ^n and Δ_{-1}^n schedules defined above. But in certain regards the earlier analysis *does* go through unchanged. We focus on these similarities with the iid problem before turning to the differences.

6.1 Equivalences between the Markov and iid cases

Perhaps the most obvious situation in which Markov and iid cases will be equivalent to one another is when we consider perturbations to the allocations at $t+1$ and (possibly) t of an agent whose allocation was not 'envied' at t. This could either be because $t+1 = 1$ (i.e., there was no prior period from the perspective of our policymaker) or because the agent's type was the highest possible at t (and thus, by our maintained focus on the 'restricted problem', was not envied). In stating this formally it is useful to define $\pi_{\Theta}^{vec}(\theta_t^n)$ as the vector of $t+1$ probabilities over Θ conditional upon a productivity draw of θ_t^n at time t (stacked in identical fashion to π_{Θ}^{vec}). Then we can state the following:

Proposition 9 ***No extra distortions in first period and at the top:*** *Suppose the solution to the relaxed problem also solves the general problem and is interior. Then for all time periods $t \geq 1$ and any reporting history θ^t whose terminal entry θ_t is the maximal element of Θ, denoted θ_t^N, the matrices $\{\Delta^{n\prime}(0)\}_{n=1}^N$ associated with the optimal allocation matrix $X_{t+1}^*(\theta^t)$ satisfy the following condition:*

$$\left(\pi_{\Theta}^{vec}\left(\theta_t^N\right)\right)' \left[\sum_{n=1}^N \gamma_n \Delta^{n\prime}(0)\right] k = 0 \tag{6.1}$$

where k is again the 2×1 vector $\begin{bmatrix} 1 \\ -1 \end{bmatrix}$ and γ_n is the nth element of any vector γ that satisfies the *equation* $\gamma' J \pi^{vec}_{\Theta}\left(\theta^N_t\right) = 0$.

Similarly in period 1 *the matrices* $\{\Delta^{n\prime}(0)\}_{n=1}^{N}$ *associated with the initial optimal allocation matrix* X_1^* *satisfy the following condition:*

$$(\pi^{vec}_{\Theta})' \left[\sum_{n=1}^{N} \gamma_n \Delta^{n\prime}(0)\right] k = 0 \tag{6.2}$$

where γ_n *is the nth element of any vector* γ *that satisfies the equation* $\gamma' J \pi^{vec}_{\Theta} = 0$ *and* π^{vec}_{Θ} *is the initial (unconditional) probability vector across productivity draws.*

The proof of these claims merely repeats the logic contained in Proposition 7, so is omitted. All we need note is that if the agent whose $t+1$ allocations are being perturbed was not envied by any other agent at time t then we do not need to concern ourselves with ensuring the perturbation is utility-neutral at t for a potential mimicker, and this concern is the only additional problem generated by a switch to Markov transition probabilities. The agent will not have been envied if he or she was of the highest possible type at t, or if $t = 0$ – and so $t+1$ is in fact the first period of the problem.

This result implies all of the 'intermediate' intratemporal optimality conditions from the iid case (that is, those associated with differential changes to utility levels across different productivity draws at $t+1$) carry over to the Markov problem for a particular subset of reporting histories. We have additionally already shown that the two 'extreme' perturbations – that is, changes along the top indifference curve and uniform utility provision, which led to the 'no distortion at the top' and generalised inverse Euler results respectively – both carry over for all histories under Markov type processes. So all that remains is to understand how the 'intermediate' perturbations are affected when agents' prior allocations *were* envied.

6.2 Differences between the Markov and iid cases

There are two important ways in which optimality requirements do change when we switch to the Markov problem. First, the dimensionality of the space within which outcomes can be perturbed to generate *intra*temporal optimality conditions is reduced by one for all agents who were envied in the previous time period. Second, and offsetting this loss of an intratemporal condition, an additional *inter*temporal condition arises, ensuring that the cost to the policymaker of preventing mimicking is spread optimally through time. We explain these points in turn.

6.2.1 Intratemporal optimality: a dimension lost

If we are considering a perturbation that applies exclusively in period $t+1$ to the allocations of agents with a common reporting history $\widehat{\theta}^t$, such that $\widehat{\theta}_t = \theta_t^n \neq \theta_t^N$ (where the latter is the maximal element of Θ), we need to make sure that this perturbation does not affect the incentive at t for truthful reporting – either for an agent whose true type is θ_t^n or for one whose true type is θ_t^{n+1} (and thus is indifferent at the conjectured optimum between reporting $\widehat{\theta}_t^{n+1}$ or $\widehat{\theta}_t^n$). This implies that the expected utility consequences of the perturbation must be zero under both the 'truth-teller's' probability measure $\pi_\Theta(\cdot|\theta_t^n)$ and the 'mimicker's' measure $\pi_\Theta\left(\cdot|\theta_t^{n+1}\right)$. In the iid case we were able at the margin to implement any linear composite of the 'basic' perturbation matrices $\{\Delta^{n\prime}(0)\}_{n=1}^N$ provided the vector of relative weights given to each, the $N \times 1$ vector γ, satisfied $\gamma' J \pi_\Theta^{vec} = 0$ for the unique probability vector π_Θ^{vec}. Recall that the nth row of the matrix J details the marginal utility consequence of the perturbation Δ^n for each type at $t+1$, so this restriction on γ ensures the net effect of the composite perturbation on expected utility is zero under the common probability measure. In general one can always find $N-1$ linearly independent γ vectors that satisfy this condition.

By the same logic, when shocks are Markov we can preserve incentive compatibility for both truth-tellers and mimickers provided we perturb outcomes at the margin according to a composite of the basic perturbations for which the weight vector γ *jointly* satisfies *two* conditions:

$$\gamma' J \pi_{\Theta}^{vec}\left(\theta_t^n\right) = \gamma' J \pi_{\Theta}^{vec}\left(\theta_t^{n+1}\right) = 0 \tag{6.3}$$

In general one can always find $N-2$ linearly independent γ vectors for which this condition is satisfied. Hence the movement to Markov probabilities has denied us the capacity to carry out intratemporal perturbations in precisely one dimension. As in the iid case, corresponding to any such γ vector will again be an equivalent ν vector that directly lists the marginal utility effects on agents at $t+1$, given by $\nu = J'\gamma$. Lemma 1 can then be easily adjusted to cover intratemporal perturbations in the Markov case:

Lemma 3 *For all time periods $t \geq 1$, all reporting histories θ^t such that $\theta_t = \theta_t^n \neq \theta_t^N$, and any vector ν that satisfies $\nu' \pi_{\Theta}^{vec}\left(\theta_t^n\right) = \nu' \pi_{\Theta}^{vec}\left(\theta_t^{n+1}\right) = 0$ it is possible to perturb the optimal allocations $X_{t+1}^*\left(\theta^t\right)$ (assumed to be interior) in a manner that will preserve the incentive compatibility constraints of the relaxed problem in all periods whilst raising the within-period utility of an agent of type θ_{t+1}^n by an amount $\nu_n \delta$ at $t+1$ for any δ satisfying $|\delta| < \varepsilon$ for some $\varepsilon > 0$ and leaving utility in all other periods constant.*

We omit to include a proof, since the logic is identical to that of Lemma 1, except that it is applied here only to the subset of within-period perturbations admissible in the Markov case. The important point is just that the specified non-marginal perturbations can be carried out whilst preserving incentive compatibility for the relaxed problem in earlier periods (even though the basic perturbations $\{\Delta^n(\delta)\}_{n=1}^N$ and $\left\{\Delta_{-1}^n(\delta)\right\}_{n=1}^N$ may no longer be admissible). Thus the *marginal* utility effects associated with them are implementable in a manner that will keep us within the constraint set of the relaxed problem, and so must come at zero marginal resource cost when the solution to the relaxed problem is known to coincide with the solution to the general problem.

Note also that the proof of Lemma 2 will go through essentially unchanged when we restrict attention to the subset of the possible utility increment vectors ν that will permit incentive compatibility to be preserved in the Markov case. This implies that the marginal cost to the policymaker of implementing any such vector at $t+1$ for all agents with a given reporting history $\widehat{\theta}^t = \theta^t$ (such that $\widehat{\theta}_t = \theta_t^n$) will again be:

$$\pi_\Theta\left(\theta^t\right)\left(\pi_\Theta^{vec}\left(\theta_t^n\right)\right)'\left[\sum_{n=1}^N \gamma_n \Delta^{n\prime}(0)\right]k$$

for γ_n defined as before. So even though weighted *pairwise* sums of the basic perturbations $\{\Delta^n(\delta)\}_{n=1}^N$ may no longer be compatible with incentive compatibility,[39] any composite perturbations that still *are* incentive compatible in the Markov case can also still have their marginal costs expressed as a linear combination of the marginal costs of these basic perturbations. Underneath all of our linear algebra remains a set of marginal movements along within-period indifference curves, and marginal compensation payments to mimickers for these.

The required intratemporal optimality conditions can now be stated formally:

Proposition 10 ***Intratemporal optimality (Markov case):*** *Suppose the solution to the relaxed problem also solves the general problem and is interior. Then for all time periods $t \geq 1$ and any reporting history θ^t such that $\theta_t = \theta_t^n \neq \theta_t^N$, the matrices $\{\Delta^{n\prime}(0)\}_{n=1}^N$ associated with the optimal allocation matrix $X_{t+1}^*\left(\theta^t\right)$ satisfy the following condition:*

$$\left(\pi_\Theta^{vec}\left(\theta_t^n\right)\right)'\left[\sum_{m=1}^N \gamma_m \Delta^{m\prime}(0)\right]k = 0 \tag{6.4}$$

where k is again the 2×1 vector $\begin{bmatrix} 1 \\ -1 \end{bmatrix}$ and γ_m is the mth element of any vector γ that satisfies the two restrictions $\gamma' J \pi_\Theta^{vec}\left(\theta_t^n\right) = 0$ and $\gamma' J \pi_\Theta^{vec}\left(\theta_t^{n+1}\right) = 0$.

The proof again follows directly from earlier arguments so is omitted here. Together with the

[39] Recall that the iid condition (5.28) is a pairwise sum of the marginal effects of the Δ^1 and Δ^n perturbations for some $n \in \{2, ..., N\}$, with weights $-\pi_\Theta\left(\theta_{t+1} \geq \theta_{t+1}^n | \theta_t\right)$ and 1 respectively.

'no distortion at the top' condition, it implies we now have $N-1$ linearly independent optimality conditions that must hold within each time period across types that share a common prior report history. The generalised inverse Euler condition gives a further condition (in all periods except the first),[40] and there are $N-1$ binding incentive compatibility constraints. Together this implies that we are one equation short of tying down the $2N$ variables that are to be determined in each period (with the exception of the first, and for any reporting history that did *not* feature the maximal element of Θ in the preceeding period). The final step in our characterisation is to provide this missing equation.

6.2.2 Intertemporal optimality: exploiting dynamic dependencies

Recall again the basic problem faced by our utilitarian policymaker. As we saw in Section 3, the first-best solution would involve all agents facing a within-period marginal income tax rate of zero, so that the marginal utility value of a unit of extra product is equal to its marginal utility cost. At the same time, the marginal utility of consumption would be equalised across agents. When types are unobservable these objectives are mutually incompatible. The ability of higher-type agents to mimic implies they would only report their types truthfully if given substantially more utility than lower types. But by raising the tax wedge on lower types – reducing their consumption and output levels simultaneously along a within-period indifference curve – one can ensure that the marginal benefits to higher types from mimicking are reduced, appealing to the intuition that we developed when presenting Figure 5.1. This in turn reduces the utility rents that higher types can extract from the policymaker – these rents being spread at the optimum across the contemporary and subsequent periods, in a manner that satisfies the inverse Euler condition. Seen in this light, the policymaker's problem is to resolve the trade-off between the provision of wasteful amounts of current and future utility to higher types, and the use of wasteful tax wedges that impede the production of lower types.

When productivity shocks are Markov there is a third alternative. Instead of reducing higher types'

[40] Again, an intertemporal resource constraint can loosely be thought of as substituting for a dynamic optimality condition in the first time period.

utility rents through tax wedges on lower types, it is possible to do it by 'twisting' the provision of utility across states in subsequent periods, so that the expected future benefits to mimickers from a given report are reduced, even whilst the expected benefits to truth-tellers are held constant. That is, if an agent were to report some $\widehat{\theta}^t$ such that $\widehat{\theta}_t = \theta_t^n \neq \theta_t^N$, it is always possible to shift allocations across states in period $t+1$ (relative to the least-cost means of providing a given level of expected utility to truth-tellers) so that agents whose true type is θ_t^{n+1} see a reduction in their expected utility from mimicking under the measure $\pi_\Theta\left(\cdot|\theta_t^{n+1}\right)$, whilst expected utility under the measure $\pi_\Theta\left(\cdot|\theta_t^n\right)$ remains unchanged. The theory of the second best suggests there will in general be net benefits to distorting $t+1$ allocations in this manner.

Before stating the main argument we must provide an equivalent to Lemma 3 to confirm incentive compatibility for dynamic perturbations. We have the following, the proof of which is in the appendix:

Lemma 4 *For all time periods $t \geq 1$, all reporting histories θ^t such that $\theta_t = \theta_t^n \neq \theta_t^N$, and any vector ν that satisfies $\nu' \pi_\Theta^{vec}\left(\theta_t^n\right) = 0$ and $\nu' \pi_\Theta^{vec}\left(\theta_t^{n+1}\right) = 1$ it is possible to perturb the optimal allocations $\left(c_t^*\left(\theta^t\right), y_t^*\left(\theta^t\right)\right)$ and $X_{t+1}^*\left(\theta^t\right)$ (assumed to be interior) in a manner that will preserve the incentive compatibility constraints of the relaxed problem in all periods whilst raising the within-period utility of an agent of type θ_{t+1}^n by an amount $\nu_n \delta$ at $t+1$ for any δ satisfying $|\delta| < \varepsilon$ for some $\varepsilon > 0$ and leaving equilibrium utility in all other periods constant.*

This result immediately takes us to the final optimality condition that we desire.

Proposition 11 ***Dynamic cost-spreading:*** *Suppose the solution to the relaxed problem also solves the general problem and is interior. Then for all time periods $t \geq 1$ and any reporting history θ^t such that $\theta_t = \theta_t^n \neq \theta_t^N$, the marginal perturbation matrices $\{\Delta^{n\prime}(0)\}_{n=1}^N$ associated with the optimal $t+1$ allocation matrix $X_{t+1}^*\left(\theta^t\right)$ together with the optimal t allocation pair $\left(c_t^*\left(\theta^t\right), y_t^*\left(\theta^t\right)\right)$ must satisfy* *the following condition:*

$$\beta R_{t+1} \frac{\tau\left(\theta_t^n\right)}{u_c\left(\widehat{\theta}_t^n ; \theta_t^{n+1}\right)\left(1-\tau\left(\theta_t^n\right)\right)+u_y\left(\widehat{\theta}_t^n ; \theta_t^{n+1}\right)} \tag{6.5}$$
$$= \left(\pi_{\Theta}^{vec}\left(\theta_t^n\right)\right)'\left[\sum_{m=1}^N \gamma_m \Delta^{m\prime}(0)\right] k$$

where k is the 2×1 vector $\begin{bmatrix} 1 \\ -1 \end{bmatrix}$ and γ_m is the mth element of any vector γ that satisfies the two restrictions $\gamma' J \pi_{\Theta}^{vec}\left(\theta_t^n\right)=0$ and $\gamma' J \pi_{\Theta}^{vec}\left(\theta_t^{n+1}\right)=-1$.

Proof. We consider a composite perturbation pair, denoted $\Delta(\delta)$ and $\Delta_{-1}(\delta)$, such that $\Delta(\delta)$ raises the within-period utility of an agent of type θ_{t+1}^m by an amount $\nu_m \delta$ at $t+1$, where ν_m is the mth entry of the vector $\nu=J'\gamma$. By earlier arguments (c.f. proof of Lemma 2), the marginal cost of this $\Delta(\delta)$ perturbation as δ is moved away from 0, assessed from the perspective of time t, will be:

$$R_{t+1}^{-1} \pi_{\Theta}\left(\theta^t\right)\left(\pi_{\Theta}^{vec}\left(\theta_t^n\right)\right)'\left[\sum_{m=1}^N \gamma_m \Delta^{m\prime}(0)\right] k$$

By Lemma 4 we know that we can remain within the constraint set of the relaxed problem through these perturbations, and the fact that the solution to the relaxed problem also solves the general problem will then imply marginal changes cannot raise a surplus. The proof of Lemma 4 shows that incentive compatibility at t is preserved by moving allocations along the indifference curve of the relevant truth-telling agent with the report history θ^t, and doing so by an amount sufficient to increase the within-period utility of a mimicker by $\beta\delta$ units. Retaining earlier definitions of the functions φ^c and φ^y, the cost of this perturbation, assessed at time t, will be:

$$\pi_{\Theta}\left(\theta^t\right)\left[\varphi^c\left(\theta_t^n, \beta\delta\right)-\varphi^y\left(\theta_t^n, \beta\delta\right)\right]$$

and so the marginal cost as δ is moved away from zero is:

$$\beta\pi_{\Theta}\left(\theta^{t}\right)\left[\varphi^{c}\left(\theta_{t}^{n},0\right)-\varphi^{y}\left(\theta_{t}^{n},0\right)\right]$$

which we have already established (c.f. proof of Proposition 6) is equal to:

$$-\beta\pi_{\Theta}\left(\theta^{t}\right)\frac{\tau\left(\theta_{t}^{n}\right)}{u_{c}\left(\widehat{\theta}_{t}^{n};\theta_{t}^{n+1}\right)\left(1-\tau\left(\theta_{t}^{n}\right)\right)+u_{y}\left(\widehat{\theta}_{t}^{n};\theta_{t}^{n+1}\right)}$$

The result then follows from the fact that the total present value of the marginal cost of the perturbation must be zero at an optimum. ■

It is well known that the shift from iid to Markov transition probabilities complicates substantially the computation of optimal dynamic policy in models such as this – the point is explored at length, for instance, by Fernandes and Phelan (2000) in the context of a dynamic agency model, and by Kapička (2010) in the context of dynamic Mirrleesian problems. Equation (6.5) provides one interpretation for why this is so: when shocks are Markov the policymaker has the capacity to spread through time the costs of any given utility advantage that mimickers have over truth-tellers, and it is always optimal to exploit this. That fact introduces an extra dynamic optimality requirement, on top of the generalised inverse Euler condition.[41] This implies one needs more information about past productivity draws when solving for an optimal within-period allocation in the Markov case than in the iid case, since one must ascertain not just the average level of the marginal cost of utility provision to implement across agent types within a period, but also the extent to which allocations should be 'twisted' to reduce prior benefits to mimicking.

It is also worth emphasising that the benefits to twisting allocations in this way are time-inconsistent.

[41] Kapička (2010) makes a similar observation when using a first-order value function method to study a specific example of a dynamic Mirrleesian model. The idea is also implicit in the general treatment of dynamic incentive provision under the first-order approach by Pavan, Segal and Toikka (2011).

As Proposition 9 shows, if $t = 1$ there would be no incentives to set a value for the object

$$\left(\pi_{\Theta}^{vec}\left(\theta_{t-1}^{n}\right)\right)'\left[\sum_{m=1}^{N}\gamma_{m}\Delta^{m\prime}\left(0\right)\right]k$$

different from zero (given $\gamma' J\pi_{\Theta}^{vec}\left(\theta_{t-1}^{n}\right) = 0$), so in all subsequent periods an 'uncommitted' policy-maker would have an incentive to revert to the least-cost means of providing a given utility distribution to agents with a known prior history.

In general, the optimality consideration highlighted here is likely to result in greater equality at $t+1$, conditional upon a given history to t, the higher is the marginal tax rate for an agent at t. This is because, as discussed, higher marginal rates are really a means for the policymaker to reduce the utility gap that has to exist between agents of adjacent types in order to prevent mimicking by the more productive. But one can also reduce this gap by reducing the benefits higher types could expect to obtain in future periods subsequent to mimicking, assessed under their type-specific probability distribution. Assuming that this latter distribution places greater weight on higher-type outcomes in the future than does the distribution specific to truth-tellers (as would be the case if $\pi_{\Theta}\left(\cdot|\theta_{t}^{n}\right)$ was first-order stochastically dominated by $\pi_{\Theta}\left(\cdot|\theta_{t}^{n+1}\right)$, for instance), one can disadvantage mimickers at t whilst leaving truth-tellers unaffected in expected utility terms by shifting $t+1$ utility away from higher types and towards lower types. Thus the 'twisting' that we have highlighted seems very likely to move outcomes towards greater equality in future *utilities* the higher are initial tax rates (and thus the greater is the distortion the policymaker is willing to accept).[42]

[42] Note that greater equality in utilities will generally require greater distortions to the production efficiency of lower types, with downward movements along their within-period indifference curves relative to the iid case, so that higher types are not given an incentive to mimic. Farhi and Werning (2010) use a simulated model with Markov transitions to show that average effective marginal income taxes across types do indeed increase through time – entirely consistent with our theoretical result.

7 Martingale convergence results

The final major area on which it is worth focusing attention is the evolution of optimal outcomes over time, and in particular at the limit as the time horizon becomes large. Suppose that the real interest rate were in all time periods equal to the inverse of the discount factor β. Then the generalised inverse Euler equation can be written as:

$$\begin{aligned} & \frac{1-\alpha\left(\theta_{t}\right)}{u_{c}\left(\theta_{t}\right)+u_{y}\left(\theta_{t}\right)\alpha\left(\theta_{t}\right)} \\ & =\sum_{\theta_{t+1}\in\Theta}\pi_{\Theta}\left(\theta_{t+1}|\theta_{t}\right)\frac{1-\alpha\left(\theta_{t+1}\right)}{u_{c}\left(\theta_{t+1}\right)+u_{y}\left(\theta_{t+1}\right)\alpha\left(\theta_{t+1}\right)} \end{aligned} \tag{7.1}$$

That is to say, we have a martingale in the marginal cost of (locally incentive compatible) utility provision. When preferences are separable between consumption and labour supply, $\alpha\left(\theta_{t}\right)=0$ holds, and the expression collapses to a martingale in the inverse of the marginal utility of consumption – an object that is strictly positive and (under the Inada conditions that we have assumed) bounded below at 0. As many authors have observed, this boundedness allows the application of Doob's martingale convergence theorem, which implies almost sure convergence in the inverse marginal utility of consumption to a finite (possibly random) limit. If one can also show that the optimum will never involve consumption staying fixed at a non-zero value (which is a likely consequence of the policy-maker's ever-present need to provide incentives[43]), convergence to zero consumption becomes the only possibility.

To generalise these results to the case at hand we need to put a bound on the object in (7.1) – the marginal cost of utility provision – for preference structures more general than the separable case. A first step is the following.

[43] In a useful discussion, Kocherlakota (2011) notes the possibility of convergence in consumption to one of the endpoints of some bounded interval of the real line in the event that the marginal disutility of labour supply is bounded away from zero and total labour supply has an upper limit. The intuition here is that when agents are sufficiently 'wealthy' or sufficiently poor they will, respectively, work zero or the maximum possible number of hours whatever their productivity draw – so stable consumption is possible following convergence to these limits.

Lemma 5 *Under an optimal plan that solves the restricted problem,* $u_c(\theta_t) + u_y(\theta_t)\alpha(\theta_t) > 0$ *always holds.*

Proof. By definition

$$\alpha(\theta_t^n) = \frac{u_c(\theta_t^n) - u_c\left(\widehat{\theta}_t^n; \theta_t^{n+1}\right)}{u_y\left(\widehat{\theta}_t^n; \theta_t^{n+1}\right) - u_y(\theta_t^n)} \tag{7.2}$$

for $n < N$, and $\alpha\left(\theta_t^N\right) = 0$. In the latter case the result follows immediately from $u_c(\theta_t) > 0$. In the former case we have from equation (2.5):

$$\frac{u_y\left(\widehat{\theta}_t^n; \theta_t^{n+1}\right)}{u_y(\theta_t^n)} < \frac{u_c\left(\widehat{\theta}_t^n; \theta_t^{n+1}\right)}{u_c(\theta_t^n)} \tag{7.3}$$

Rewriting our object of interest, we have:

$$u_c(\theta_t) + u_y(\theta_t)\alpha(\theta_t) = \frac{u_c\left(\widehat{\theta}_t^n; \theta_t^{n+1}\right) - u_c(\theta_t^n)\frac{u_y\left(\widehat{\theta}_t^n; \theta_t^{n+1}\right)}{u_y(\theta_t^n)}}{1 - \frac{u_y\left(\widehat{\theta}_t^n; \theta_t^{n+1}\right)}{u_y(\theta_t^n)}} \tag{7.4}$$

The numerator of the right-hand side is clearly positive by the preceeding inequality, and the denominator likewise by the fact the marginal disutility of production is lower for higher types (c.f. inequality (2.2)). ■

Given the definition of $\alpha(\theta_t)$ this allows us almost immediately to state a bound when consumption and labour supply are Edgeworth substitutes. But when they are Edgeworth complements our scope for doing so proves surprisingly limited. Taken together we have the following result.

Lemma 6 $\frac{1-\alpha(\theta_t)}{u_c(\theta_t)+u_y(\theta_t)\alpha(\theta_t)} > 0$ *always holds under an optimal plan that solves the restricted problem, unless (a) consumption and labour supply are Edgeworth complements, and (b) productivities follow a non-iid process.*

Proof. With separability between consumption and labour supply $\alpha(\theta_t) = 0$, and the assumption

$u_c(\theta_t) > 0$ is enough to confirm the result. When consumption and labour supply are Edgeworth substitutes we have $\alpha(\theta_t^n) < 0$ (the marginal utility of consumption is higher for mimickers than truth-tellers, since the former need not work so hard to produce a given level of output), and the result follows from Lemma 5. When consumption and labour supply are Edgeworth complements it is possible to prove the bound only for the iid case. The reasoning is far more involved, and we relegate it to an appendix. ■

Having put a zero lower bound on the marginal cost of utility provision for these specific cases, when $R_t = \beta^{-1}$ for all t a direct application of Doob's martingale convergence theorem implies the object $\frac{1-\alpha(\theta_t)}{u_c(\theta_t)+u_y(\theta_t)\alpha(\theta_t)}$ must converge almost surely along all realisations of θ^∞ to some value $X \in [0,\infty)$, where X is potentially a random variable. We want to be able to say more about the value of X. In fact, it turns out – as in the separable case – that X must equal zero. The next Proposition establishes this.

Proposition 12 ***Convergence:*** *Suppose $R_t = \beta^{-1}$ for all $t \geq 1$. Then $\frac{1-\alpha(\theta_t)}{u_c(\theta_t)+u_y(\theta_t)\alpha(\theta_t)} \overset{a.s.}{\to} 0$ holds under any optimal plan that solves the restricted problem (and is interior at all finite horizons), unless (a) consumption and labour supply are Edgeworth complements, and (b) productivities follow a non-iid process.*

Proof. See appendix. ■

This result is an obvious generalisation of the 'immiseration' results obtained by studying convergence of the standard inverse Euler condition. Moreover, almost sure *consumption* immiseration (in the sense that inverse of the marginal utility of consumption – and hence consumption itself – must tend to zero for almost all agents) is a direct implication of this result, when one recalls that $\frac{1-\alpha(\theta_t)}{u_c(\theta_t)+u_y(\theta_t)\alpha(\theta_t)} = \frac{1}{u_c(\theta_t)}$ when $\theta_t = \theta_t^N$ (the highest type): the outcome for an agent who draws the top productivity parameter in the tth period *must* be immiseration (almost surely) at the limit as t becomes large, and incentive compatibility then demands that all lower types with the same history must have a still worse lot. So the more complicated nature of the expression for the marginal cost of utility provision in the non-separable case does not undermine the extreme predictions regarding long-run consumption when martingale convergence *can* be applied. The political difficulties associated with long-run commitment to a scheme with such severe future outcomes are plainly immense, even abstracting from the more fundamental question of whether the welfare of the initial period's cohort of agents *ought* to be the exclusive concern for public policy.[44] For this reason alone the immiseration result is a troubling one: it is hard to imagine a scheme more likely to result in government default than one that demands its future citizens should be enslaved to pay the debts of the past.[45]

Perhaps the more surprising result of this section, though, is that when productivity follows a Markov process and consumption and labour supply are Edgeworth complements – so that those who are working longer hours with a given level of consumption have a higher marginal utility of consumption – we *cannot* put a zero lower bound on the marginal cost of utility provision. Indeed, it is quite possible that this marginal cost may turn negative. This possibility we are able to confirm through a finite-horizon computed example, the details of which we now present.

7.1 Computed example

We assume that production is linear in labour supply, with the marginal product of labor equal to θ, and that the utility function takes the form outlined in King, Plosser and Rebelo (1988):

$$u(c, y; \theta) = \frac{c^{1-\varsigma}}{1-\varsigma} \exp\left\{(\varsigma - 1) v\left(\frac{y}{\theta}\right)\right\} \tag{7.5}$$

[44] This latter question is explored in detail by Farhi and Werning (2007).

[45] Clearly if consumption is reaching zero at the limit then the within-period surplus raised for almost all histories must be substantial. We know, for instance, that 'top' agents will certainly be producing very large quantities of output, since $u_c + u_y = 0$ for these types. This surplus must be being used either to service interest on outstanding debts or to fund the lavish consumption of some measure-zero subset of agents whose luck has never been out. The latter is probably even less polictically plausible than the former.

with the labour disutility schedule v defined by:

$$v(l) = \frac{l^{1+\upsilon}}{1+\upsilon} \tag{7.6}$$

This function implies that consumption and labour supply are Edgeworth complements provided $\varsigma > 1$, and are Edgeworth substitues for $\varsigma < 1$.

A substantial practical advantage of the solution method presented in this paper is that it provides a complete set of equations necessary to solve any given example – so provided there is a finite number of types and of time periods, for any given parameterisation we can obtain a solution simply by solving these equations numerically. Specifically, if T is the total number of time periods and N the cardinality of Θ then we will have $\sum_{t=1}^{T} 2N^t$ variables to tie down in total (in each period, an output and consumption level for an agent of each current type, for each history). The method presented above delivers precisely this number of equations, which can be jointly solved to machine accuracy using standard non-linear solution algorithms. Unlike methods that exploit value function iteration, the approach is equally fast whether shocks follow an iid or a Markov process, with the latter simply involving a slightly different set of equations.

For our example we assume two types (identical across all time periods): θ_L and θ_H, with $\theta_L < \theta_H$. Transition probabilities are denoted as follows:

$$\begin{aligned}
\pi_\Theta(\theta_t = \theta_H) &= P^H & if \quad t = 1 \\
\pi_\Theta(\theta_t = \theta_H | \theta_{t-1} = \theta_H|) &= P_H^H & if \quad t > 1 \\
\pi_\Theta(\theta_t = \theta_H | \theta_{t-1} = \theta_L|) &= P_L^H & if \quad t > 1
\end{aligned}$$

We set $T = 6$, implying 252 variables to determine. Since at this stage the purpose of the example is more to find a counterexample to $\frac{1-\alpha(\theta_t)}{u_c(\theta_t)+u_y(\theta_t)\alpha(\theta_t)} > 0$ than to claim realism *per se*, and since this counterexample is more likely to arise in our finite horizon the greater is the value of ς,[46] we choose

[46] High values of ς imply strong complementarity, and thus a much lower marginal utility of consumption for mimickers at

the relatively high value: $\varsigma = 10$. For the other parameters we choose values $v = 2$ and $\beta = 0.99$. We normalise $\theta_L = 1$ and set $\theta_H = 2$. The inital probability P^H we set to 0.5, with strong type persistence thereafter: $P_H^H = 0.9$ and $P_L^H = 0.1$.

Figure 7.1 is a histogram summarising the distribution of the marginal cost of utility provision across agents in the 6th (and final) period of the simulation, with bins 0.1 units wide (the units here being the single consumption good). The high degree of persistence accounts for this distribution's clear bimodal character.[47] What is of more interest is that the marginal cost of utility provision (provision, that is, in a manner that preserves *within*-period incentive compatibility) is negative for exactly half of the agents in this period. These agents are the half of the population with contemporaneous productivity θ_L.[48]

A negative marginal cost of utility provision also obtains for almost all low-type agents in the 5th period of the simulation, so the result is not dependent upon the period in question being the last. On the surface it is a very counter-intuitive outcome (surely the policymaker can provide utility to a subset of agents and generate a surplus?), so it is worth examining it in detail. Recall that when consumption and labour supply are Edgeworth complements, a provision of utility by consumption increments alone at a given output level would benefit low types by more than (mimicking) high types, since the latter supply less labour to produce the given quantity of output – and thus do not benefit from complementarities to so great an extent. Hence to preserve incentive compatibility (for utility movements in either direction) any consumption increment must be accompanied at the margin by an increase in production, which causes greater marginal disutility to lower types than higher (the former are already working longer hours, so their marginal disutility of effort is greater), eliminating the utility imbalance.

a given allocation than for truth-tellers. To offset this requires utility provision along a vector that will increase production requirements significantly alongside any extra consumption provision (this exploits the higher marginal disutility of production on the part of truthtellers). Since the marginal cost of utility provision is lower the more output is increased for a given consumption increase, higher complementarity is likely to be associated in general with lower marginal costs.

[47] Roughly three fifths of agents draw the same type in all six periods.

[48] Recall that high-type agents must be associated with positive marginal costs, since $\alpha(\theta_H) = 0$ always holds.

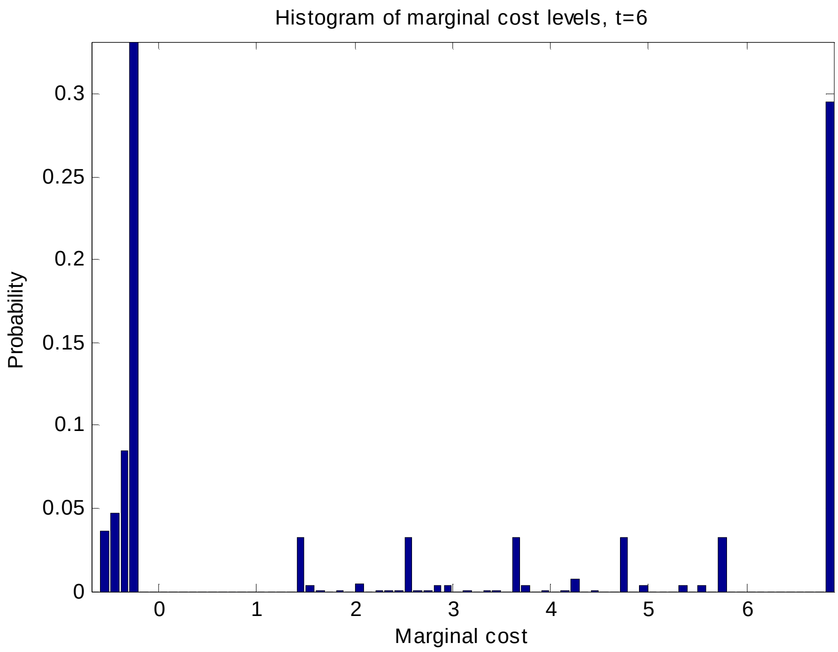

Figure 7.1: Distribution of marginal cost of utility provision in 6th period

The results of the simulation suggest that the choices of low types are, at the optimum, being distorted sufficiently far away from a point at which the slope of their within-period indifference curve equals one that even movement along a vector giving *equal* consumption and output increments would still raise their utility by more than it would raise the utility of high-type mimickers – and so output must be increased by *more* than consumption at the margin to obtain balance. Notice that this suggests the output of low-type agents is being restricted substantially at the optimum: the lower is output the lower is the *difference* in the marginal effect on utility of an increase in it between truth-tellers and mimickers, and so the more it must be raised for an incentive-compatibility-preserving perturbation.

Why is it not possible to exploit the negative cost of utility provision to generate a surplus? It is simply that there does not exist a means to provide utility to a given agent in a way that generates resources *whilst at the same time offsetting any effects on incentive compatibility constraints.* A gift of extra utility to a low-type agent in the 6th period would induce high-type agents with the relevant prior history to switch to a mimicking strategy. The cost of preventing this, through an equal utility increment to a high-type agent, may directly offset the generation of a surplus. Even if not, incentive compatibility in the 5th period would also be violated if we are considering allocations to those whose prior report was $\widehat{\theta}_L$. Equally in the 5th period, a gift of utility to a low-type agent whose marginal cost is negative could be incentive-compatible if accompanied by a reduction in utility across all agents in the subsequent period; but the aggregated present value of the (negative) costs of these perturbations will be zero, by the generalised inverse Euler condition. Ultimately, no matter what composite marginal perturbation one tries to construct, either local incentive compatibility must be violated, or no surplus raised.

The important question that follows from these results is whether the *potential* for a more benign long-run outcome than immiseration is indeed likely to be realised in the event of complementarity: just because we cannot prove it by martingale convergence does not mean immiseration can be ruled out. One can only conjecture in the absence of a full solution to the infinite-horizon model, but there are reasonable economic grounds for believing immiseration will be avoided. Specifically, note that the tendency towards immiseration (when it does hold) must derive in part from the finite stock of resources at the policymaker's disposal. As time progresses, either a prior tendency to front-load utility provision through debt finance, or promises of very high utility levels to a measure-zero (perpetually lucky) subset of agents, or some combination of the two, results in the maximum possible surplus being extracted from almost all agents. But if in the case of complementarities the marginal cost of raising the utility of agents turns negative then a tendency to immiserate may well be counter-productive – costing resources rather than generating them. Clearly the policymaker has no *direct* desire to see immiseration occur, so it seems unlikely that this cost will be worth paying.

7.2 Linking saving wedges and immiseration

The results of this Section – in particular Lemma 6 – allow for a slight extension to the set of circumstances in which we can claim it is optimal to deter savings (in some meaningful sense). We can state the following.

Proposition 13 ***Deterred savings (3):*** *Suppose the solution to the relaxed problem also solves the general problem and is interior. Then for all time periods $t \geq 1$ and for all reporting histories θ^t, if consumption and labour supply are Edgeworth complements then savings will be deterred at the optimum, in the sense that the allocations $\left(c_t^*\left(\theta^t\right), y_t^*\left(\theta^t\right)\right)$ and $X_{t+1}^*\left(\theta^t\right)$ will satisfy inequality (5.23), with that inequality holding strictly so long as the object $\frac{u_c(\theta_{t+1})+u_y(\theta_{t+1})\alpha(\theta_{t+1})}{1-\alpha(\theta_{t+1})}$ varies for different draws of $\theta_{t+1} \in \Theta$.*

The proof is identical to that of Proposition 5, which can be applied whenever the bound:

$$\frac{1-\alpha\left(\theta_t\right)}{u_c\left(\theta_t\right)+u_y\left(\theta_t\right)\alpha\left(\theta_t\right)} > 0 \tag{7.7}$$

holds – which we now know to be the case under complementarity and iid productivity draws by Lemma 6. What is interesting here is that the cases in which we can say with certainty that it is optimal to deter savings (relative to some optimality criterion that would have to hold under autarky) are precisely the cases in which we can confirm immiseration as a limiting outcome: essentially, all situations *except* that of Markov productivity draws and complementarity. This is unlikely to be a coincidence. If savings are being distorted at the optimum, the policymaker is implicity choosing to 'front-load utility' in expectation. This is just a direct reading of inequality (5.23). But if utility is being front-loaded it would not be at all surprising if the policymaker's wealth were deteriorating continually over time – so that outstanding debt obligations become cripplingly large as time passes. In this case agents in the economy would have to put in large amounts of work for little or (at the limit) no return, just to preserve the tax scheme's solvency. This implies immiseration. Only when the optimality of 'front-loading' utility no longer necessarily goes through can we escape this 'trap'.

8 Conclusion

The main contribution of this paper is a methodological one. Dynamic models with asymmetric information are a growing source of interest to macroeconomists, and the dynamic version of the Mirrlees income tax problem has generated particular interest. But practically all of the analysis of these models to date has relied on the recursive computation of value functions, defined by a Bellman-type operator appropriately augmented to ensure past promises are kept. These methods are extremely powerful and widely applicable, but their results can be difficult to interpret, simply because it is not always clear exactly which trade-offs have contributed to generating a given policy function or time-path for a variable of interest. Our analysis gives an alternative means to gain insight into this class of problems, through carefully-chosen perturbations to optimal allocations. In particular, we appeal to the revelation principle to treat the optimum as one in which individuals make direct reports of their types, and investigate how to perturb allocations along dimensions chosen to ensure there will be no changes to these reports – at least for small enough perturbations. This approach allows us to obtain a complete set of optimality conditions that, together with the binding incentive compatibility restrictions and an aggregate resource constraint, are sufficient to characterise the problem's solution. The method is analogous to solving consumer choice problems by asserting that marginal rates of substitution must equal price ratios: in our case as there, the relevant optimality conditions do not *directly* make use of the problem's constraints, relying instead on perturbations that, by construction, will hold constant any values that feature in them (expenditure in the case of consumer choice, relative returns from mimicking here).

An important limitation to our approach is that we must know in advance exactly which incentive compatibility constraints bind at the optimum. In the static Mirrlees problem the single crossing condition is known to ensure these constraints bind 'downwards' locally, and we present sufficient conditions relating to the optimal allocation that can be checked to verify whether this extends to the dynamic case for any given example. We proceed under the assumption that it does, but requiring *ex ante* knowledge of this essential characteristic of the solution is undoubtedly a disadvantage. Developing the sufficiency conditions into a more easily interpretable form, particularly when shocks are non-iid, is an important area for future research.

The optimality conditions that we derive are easiest to understand through a graphical representation of the problem in output-consumption space. They are a set of cross-restrictions on (a) the cost to the policymaker of moving 'along' each agent's agent's within-period indifference curve, reducing that agent's consumption and output jointly, and (b) the cost of providing a unit of utility to each agent in such a way that a mimicking higher-type agent would receive the same utility increment. Appropriately-chosen composites of these movements, either within or across periods, can ensure local incentive compatibility always continues to hold, and so cannot be applied in the neighbourhood of the optimum in a way that would generate a surplus for the policymaker.

This analytical method is likely to be very useful from a computational perspective, since it eliminates any need to solve maximisation problems directly when calculating the optimum to a given problem. Instead, one need only impose (jointly) the complete set of equations known to characterise that optimum. When the problem has a finite and sufficiently small number of time periods, and relatively small set of productivity types, the solution can be established to machine accuracy by solving a quite manageable set of simultaneous equations. In an infinite horizon problem functional approximation will still be necessary, since future values feature in incentive compatibility constraints, but these values should be expressible as functions of a relatively small set of variables, and will not have to be defined by any supremum operator (or similar). In the iid case, for instance, intertemporal optimality can be ensured by linking outcomes at $t+1$ for agents with a common history simply to the marginal cost of utility provision to those agents at t. This marginal cost variable alone should then be enough to establish the value function.

But the focus of the paper has been on exploiting the analytical results that a perturbation approach can expose, and here there are several. On a higher theoretical level, we have shown that when productivity draws are iid the problem separates into intratemporal and intertemporal dimensions, with the set of intratemporal optimality restrictions that must hold being identical to those that are necessary in a static model, and a single dynamic optimality condition all that is required to ensure an optimal use of resources through time. In the more realistic case that productivity draws follow a Markov process with persistence, one extra dynamic optimality condition emerges – reflecting an extra ability that the policymaker now has to exploit differences in productivity measures between mimickers and truth-tellers, in order to spread distortions through time. Accompanying this is a reduction by one in the number of intratemporal optimality conditions that can be stated. Rather like the use of separability in utility functions to simplify the statement of optimal consumption choices, this partition of the problem can, it is hoped, make the character of its solution much easier to understand.

From a more practical perspective, we have shown that many of the well-known results from static income tax theory generalise to the dynamic case. In particular, regardless of whether the shock process is Markov or iid we can show that effective within-period marginal income tax rates are always weakly positive at the optimum – in the sense that the solution always involves individuals being willing to produce at the margin for a return that is (weakly) less than their marginal product. Moreover, agents whose type is the highest always have a zero effective marginal tax rate, and these are the only agents who do so.

Turning to savings taxes, it is already well-known that in the event of separability between consumption and labour supply it is optimal to apply a positive tax wedge to savings, in the sense that the marginal utility of consumption in period t is below its expected value at $t+1$ (allowing for discounting and the interest rate): this follows from the well-known 'inverse Euler equation' that holds in that case, combined with Jensen's inequality. We have been able to generalise this result in two regards. First, and rather limited in its scope, we have shown that the marginal utility of consumption for an agent whose productivity type is the highest possible must also be below the value it would take under autarky (relative to its expected value in the following period) when consumption and labour supply are Edgeworth substitutes. But one need not focus simply on the *consumption* Euler equation as characterising dynamic optimality: the marginal rate of substitution between output levels in one period and the next, or between arbitrary vector combinations of consumption and output in one period and the next, must likewise equal the intertemporal price ratio at any autarkic allocation. Specifically, the inverse of the marginal cost of incentive-compatible utility provision is the marginal utility associated with a particular joint change in consumption and output, and the existence of an optimality condition relating to this object allows us to confirm that savings are always deterred at the optimum (in an economically meaningful sense) unless consumption and labour supply are Edgeworth complements *and* productivity draws are non-iid.

This latter result has strong connections with the final area that we have investigated in detail: allocations in the long run. Once again, except in the case that consumption and labour supply are Edgeworth complements and productivity draws are Markov, we have been able to put a zero lower bound on the marginal cost of incentive-compatible utility provision – which in turn will follow a martingale process in the event that the real interest rate equals the inverse of the discount factor β. Martingale convergence theorems then imply almost sure immiseration for all agents in the economy under standard preference assumptions. With complementarity and Markov shocks we have shown by counterexample that the marginal cost of utility provision can in fact turn negative, and so immiseration need not take place. Indirectly this result seems to shed some light on the cause of immiseration under alternative assumptions: the fact that it need not occur in precisely the same case that savings need not be deterred at the optimum suggests a connection between the implicit decision on the part of the policymaker to front-load the provision of utility when savings *are* being deterred – a strategy that is likely to involve some initial borrowing – and immiseration as the costs of servicing the resulting public debt accumulate.

Finally, we note that the methods used in this paper can be applied more widely, albeit with some adaptation. For instance, a companion paper outlines a similar perturbation method applicable to dynamic Mirrlees problems in which the type space Θ is a continuum. Reassuringly all of the results from this paper extend to that case in the natural way, and we are able to provide an expression for optimal marginal tax rates in the *static* model that is considerably simpler than those available in the literature to date. A second area of applicability is to dynamic agency models, where a similar set of optimality conditions can be derived under the assumption that the 'first order approach' is valid.

Part III

Optimal time-invariant policy from behind a veil of ignorance

1 Introduction

What are the appropriate considerations for the design of long-term macroeconomic policy? It has been accepted for some time now that certain aspects of policy choice – notably the determination of monetary policy – should be beyond the day-to-day control of democratically elected executives, with independent institutions or pre-determined legislation instead governing choice variables under normal circumstances. This has been influenced heavily by a belief that the associated policy questions may be affected by a time-inconsistency problem – a lack of coherence between the preference orderings of policymakers who exist at different points in time over economic variables that are set in one *particular* time period. This problem is generally considered to follow from the impossibility, *ex post*, of influencing prior expectations that it was desirable to *plan* to influence in a certain manner *ex ante*.[49]

But despite the global consensus on the merits of central bank independence, in particular, that has emerged over the past two decades, there remain surprisingly few attempts at a clear general theory regarding the appropriate objectives to ask independent institutions to pursue, or (more generally) the appropriate normative principles to enshrine in lasting economic policy legislation when time-inconsistency matters. For instance, the famous work of Rogoff (1985) studied the very specific problem of finding the optimal degree of central bank 'conservatism' in an inflation stabilisation model in the

[49] The literature of the early 1990s on the desirability of central bank independence was heavily influenced by this understanding of the monetary policy problem. For instance, in an influential paper Alesina and Summers (1993) claimed that "dynamic inconsistency theories of inflation ... make it plausible that more independent central banks will reduce the rate of inflation."

The decisions to grant operational independence to the Bank of England in 1997 and to give complete freedom even over policy objectives to the new European Central Bank were both functions of this consensus view.

spirit of Barro and Gordon (1983), whilst the influential proposal by Walsh (1995) for optimal central bank contracts was similarly particular to a given linear-quadratic policy problem. What is lacking is a set of universal principles for resolving disagreement between policymakers situated at different points in time, *involving neither the elevation of the priorities of one policymaker above all others, nor resignation to a non-co-operative outcome* (in the form of a Stackelberg, or 'discretionary' equilibrium that is not particularly desirable to any policymaker).

The aim of this Part of the thesis is to present, and explore the properties of, a new method for determining choice in environments subject to this form of time-inconsistency – a method that we argue resolves the disagreements between policymakers in different time periods in as disinterested a manner possible. Specifically, we want to ask what are the consequences for optimal policy of assuming that choices must be taken *without knowledge of the time period in which they are to be implemented.* We see this as the most direct way to eliminate the source of conflict among the policymakers who exist in different periods, without imposing artificial adjustments to their respective objectives. Instead, it is effected through a series of transformations to the general problem that are best understood as restrictions on the policymaker's information set.

We justify this 'restricted information' approach by appeal to the extensive literature in the social contractarian political tradition, which argues that institutional design *should* be conducted in a disinterested fashion – neglecting priorities that are particular to the circumstances of any subset of the designers. The most famous modern articulation of this position was that of Rawls (1971), and although his principal focus was on institutions to mediate interpersonal rather than intertemporal conflicts of interest, the idea that the institutions in which we are interested might be chosen from an 'original position' in which "it should be impossible to tailor principles to the circumstances of one's own case" can be applied fruitfully to the time-inconsistency problem – achieving a coherent, desirable solution without imposing the tyranny of any one policymaker's perspective. The transformations to the main choice problem that we apply can, accordingly, be likened to the application of a 'veil of ignorance' from behind which it is impossible to infer anything about the time period for which choice is being made.

Rawls's work focuses in particular on the normative idea of *justice*. It is for this reason that it has special relevance to the problems we study. When time inconsistency applies to an economic problem it is vital that any long-run commitment strategy, fixed by legislation or institutional devolution, is seen to be 'just'. If not, future generations will always be free to change the appropriate legislation as they see fit, and are unlikely to be prevented from doing so by social norms that limit their pursuit of short-term self interest – given that the original 'unjust' legislation was presumably itself tailored to the specific interests of a *prior* generation. A 'just' settlement is far more likely to be a robust one, as any challenge to it would pit short-term expediency against that higher principle. This is the *practical* motivation for the policy approach we develop: in many areas of economic interest, social norms are the only commitment devices available.

Denying the policymaker knowledge of time is a more complicated and subtle matter than it may first appear. In particular, we do not wish to deny the possibility of *state*-contingent policymaking. If there has been a positive cost-push shock, for instance, one would certainly want monetary policy to respond to this. But we *do* want to avoid the possibility that knowledge of these states could be exploited to *infer* the time period (with some degree of precision), and thus to pursue a policy whose optimality depends on the irrelevance of past expectations to current outcomes. Most of the analysis that follows is focused on the problem of presenting the policymaker with an information set to which the time period is orthogonal, whilst keeping this set sufficiently rich that the set of policy options remains suitably large.

The analysis that we develop highlights an interesting and intuitive asymmetry. When choosing variables that feature in expectational constraints, we show that a policymaker denied all knowledge of time should act in a manner that also maximises the (expected) 'steady state' value of welfare. This will differ from conventional 'commitment' policy through the irrelevance of the time preference discount factor to optimal choice. But when choosing 'state' variables (an asset stock, quantity of capital, or similar), whose lagged values feature as structural constraints on outcomes, all of the standard results under discounting instead apply; the rate of time preference is then an important parameter once more. In this sense our approach provides a natural link between the application of standard ('time-0') optimisation in models where no expectational constraints feature – and thus time-inconsistency is not an issue – and of the 'unconditional' optimisation approach set out by Taylor (1979), and recently analysed by Damjanovic, Damjanovic and Nolan (2008), when dealing with 'forward-looking' problems.

The advantages to taking a 'veil of ignorance' approach to optimality are best seen by examining its consequences in specific models. Such examples should also clarify exactly why certain steps are taken in the general theoretical presentation that may appear opaque when viewed in isolation. But it is necessary to provide the general analysis first – if only so that a 'veil of ignorance' policy can be well defined for the examples. This general analysis therefore follows in Section 2, whilst Sections 3, 4 and 5 apply the theory to models of inflation bias, redistribution under participation constraints, and optimal dynamic tax policy respectively. The notation in Section 2 soon becomes quite involved – a consequence of our desire to link choice to a potentially rich subset of the *history* of variables, but not the time period in which they obtain. In an attempt to ease the burden of this on the reader we build the analysis up in stages, first presenting the relevant arguments for models in which there are no endogenous state variables, before moving on to the considerably more complicated setup in which these objects do feature. Since the example of Section 3 has no endogenous states, the reader may find it preferable to cover this immediately after the analysis of the simpler case in Section 2, before tackling the rest of the general analysis.

2 Analysis

2.1 Basic framework

We first set out the general structure of the policy problems that we wish to consider, nesting numerous specicfic examples of time-inconsistent environments.

A policymaker operating at time 0 has preferences over dynamic economic outcomes described by W_0 (with higher values of W_0 preferred): the expected value of the discounted sum of a time-invariant within-period objective function, given by:

$$W_0 = \sum_{t=0}^{\infty} \beta^t \int_{H^\infty} \pi \left(x\left(h,t\right),\varepsilon\right) dF_t \left(h|h^0\right) \tag{2.1}$$

β here is the policymaker's discount factor. $x : H^\infty \times \mathbb{Z} \rightarrow X$ is a mapping determining the t-dated realisation of the vector of endogenous variables $x \in X$ from the infinite history of all endogenous and exogenous variables $h \in H^\infty$ (where this history includes contemporary *exogenous* variables, but only lagged values of endogenous variables) together with the time period $t \in \mathbb{Z}$ (where $\mathbb{Z}$ is the set of integers, including zero for convenience); and $\varepsilon \in E$ is the contemporary vector of iid exogenous variables. The infinite history of exogenous vaiables is denoted $h_\varepsilon \in E^\infty$,[50] with $(E^\infty, \Sigma_\varepsilon^\infty, F)$ describing an exogenously fixed probability measure F on Borel subsets of E^∞, $\Sigma_\varepsilon^\infty$. The mapping $x\left(h,s\right)$ together with this exogenous probability measure and an initial history h^0 will then allow a time-specific conditional probability measure to be defined on the complete history space, with associated triple $\left(H^\infty, \Sigma^\infty, F_{t|H^0}\right)$. $F_t\left(A|h^0\right)$ is then used to denote the probability a history in the set $A \subseteq H^\infty$ will characterise time t, conditional on observing the specific initial history h^0. Probability measures conditional on subsequent histories (realisations up to some time period $s > 0$) are defined in a similar fashion. Note that this distribution depends upon the policy chosen in each time period from 0 onwards. E and X are both assumed to be subsets of $\mathbb{R}^n$ for some n. For simplicity, the within-period

[50] So note that we are using H^∞ as shorthand for $E^\infty \times X^\infty$.

objective function $\pi(\cdot)$ is assumed to be strictly quasi-concave in x for all ε, and continuously differentiable in all of its arguments. (The latter assumption is not essential to the arguments that follow, though it aids exposition.)

The policy choice is across different values of the function $x(h,s)$, subject to two sets of constraints. First, the 'forward-looking' restrictions:

$$\sum_{s=0}^{T} \int_{H^{\infty}} g_s\left(x\left(h, t+s\right), \varepsilon\right) dF_s\left(h|h'\right) \geq 0, \tag{2.2}$$

for all $h' \in H^{\infty}$ and all $t \geq 0$, where $g_s(\cdot)$ may be vector-valued in the event that more than one constraint features. In general we assume that there are N such constraints. Note that we are assuming only contemporary ($t+s$-dated) endogenous variables feature in the constraint function $g_s(\cdot)$. This additive separability could probably be relaxed but only at substantial notational cost, and we choose not to pursue this. We use T to denote the maximum horizon at which future variables influence current constraints: T may be infinite in some models of interest, particularly when infinite-horizon (dynastic) utility levels enter the constraint set. We assume that each $g_s(\cdot)$ function is weakly concave in x, so that the constraint set is convex for any given ε, and also that it is continuously differentiable (again, the latter assumption is made chiefly for expositional purposes). One simple example of a forward-looking constraint is the linear New Keynesian Phillips Curve, which only features the one-period-ahead expectation of inflation (so $T=1$). Another example is a constraint that a household must be promised a utility level sufficiently great to ensure their participation in a social insurance scheme from which they may opt out.

In addition there are M 'backward-looking' constraints, jointly taking the form:

$$m\left(x\left(h,t\right), z\left(h\backslash 1, t-1\right), \varepsilon\right) \geq 0, \tag{2.3}$$

which must likewise hold for all histories and all t, where $z \in Z$ is a subset of the choice variables, *defined* as the set that features with a lag in constraint (2.3). These are the model's 'endogenous state variables'. In what follows we denote the remaining 'control variables' in x by $y \in Y$, so that $X = Y \times Z$. The notation $h \backslash 1$ is used to denote the history obtained by deleting most recent period's entries in h. Again, we assume that $m(\cdot)$ is weakly concave in (x, z) in order to keep the constraint set convex, and that it is continuously differentiable. We can assume without loss of generality that if a particular variable features as an argument of a function $g_s(\cdot)$ for any $s > 0$ then it is not contained in z – defining new contemporaneous 'auxiliary' variables if necessary to eliminate lagged endogenous variables from constraints that feature forward-looking terms.[51]

2.2 Time-inconsistency: existing approaches

Without further restrictions the solution to the problem of maximising W_0 subject to (2.2) and (2.3) will be time-inconsistent, in the sense that the optimal rules devised for the vector of endogenous variables $x(h, t)$ will vary depending on the time period in which optimisation takes place, holding fixed the history h. Here we show why this problem arises – in terms of the notation that we have adopted – and set out the main approaches that exist in the literature to resolving the difference in opinion between policymakers at different points in time.

The time-inconsistent character of optimal plans can be seen most easily by inspecting our general problem's first-order conditions, taken under the assumption of an interior solution and using Lagrange multipliers to incorporate the constraints. Denoting by $\lambda_g(h, t)$ the vector of (present-value) multipliers associated with constraint (2.2) at time t after history h, and $\lambda_m(h, t)$ the corresponding multipliers on constraint (2.3), a necessary condition for maximising W_0 subject to (2.2) and (2.3) is that the

[51] This procedure is common, for instance, in the numerical analysis of DSGE models.

following should hold for almost all feasible histories h at each time period $t \geq 0$:

$$\begin{aligned} & \frac{\partial \pi\left(x\left(h,t\right),\varepsilon\right)}{\partial y_{i}\left(h,t\right)}+\sum_{s=0}^{\min\{t,T\}}\beta^{-s}\lambda_{g}\left(h\backslash s,t-s\right)^{\prime}\frac{\partial g_{s}\left(x\left(h,t\right),\varepsilon\right)}{\partial y_{i}\left(h,t\right)} \\ & +\lambda_{m}\left(h,t\right)^{\prime}\frac{\partial m\left(x\left(h,t\right),z\left(h\backslash 1,t-1\right),\varepsilon\right)}{\partial y_{i}\left(h,t\right)} \\ & =0, \end{aligned} \tag{2.4}$$

for each variable y_i in the vector of t-dated control variables, and:

$$\begin{aligned} & \frac{\partial \pi\left(x\left(h,t\right),\varepsilon\right)}{\partial z_{i}\left(h,t\right)}+\lambda_{g}\left(h,t\right)^{\prime}\frac{\partial g_{0}\left(x\left(h,t\right),\varepsilon\right)}{\partial z_{i}\left(h,t\right)} \\ & +\lambda_{m}\left(h,t\right)^{\prime}\frac{\partial m\left(x\left(h,t\right),z\left(h\backslash 1,t-1\right),\varepsilon\right)}{\partial z_{i}\left(h,t\right)} \\ & +\beta\int_{H^{\infty}}\lambda_{m}\left(h^{\prime},t+1\right)^{\prime}\frac{\partial m\left(x\left(h^{\prime},t+1\right),z\left(h,t\right),\varepsilon^{\prime}\right)}{\partial z_{i}\left(h,t\right)}dF_{1}\left(h^{\prime}|h\right) \\ & =0, \end{aligned} \tag{2.5}$$

for each variable z_i in the vector of t-dated control variables.[52]

The latter of these conditions takes an identical form irrespective of the time period, but optimality with respect to 'forward-looking' variables – those entries in y that feature in $g_s(\cdot)$ functions for $s > 0$ – demands different considerations depending on the value of t. Specifically, if $T > 0$ then the upper limit of the summation in (2.4) will depend upon t for at least one y_i – in the sense that the sum itself will feature a larger number of terms the higher is t (at least as one moves from $t = 0$ to $t = 1$). This implies that the policymaker's optimal choice of endogenous variables in time t following history h will in general differ from their optimal choice in time 0 following history h – which in turn implies that a different optimal plan would follow if the policymaker were to re-optimise at time t. This is the familiar time-inconsistency result first highlighted formally by Kydland and Prescott (1977).

[52] Recall that the vector x stacks y and z variables together.

Beyond imposing direct constraints that limit the choice set of future policymakers to a singleton, there are two methods generally used to accommodate the possibility that future policymakers re-optimise. The first, following arguments made by Kydland and Prescott, is to assume that no institutional constraints can ever meaningfully affect the choice of future policymakers, and thus that concerns for the impact of present choices on past expectations can never directly affect outcomes. In our setup, this implies that one never has the scope to impose condition (2.4) for $t > 0$, since the terms in the summation for $s > 0$ will not be relevant considerations to the policymakers who actually choose variables in periods 1 and beyond. Instead, the policymaker in each period must do the best he or she can, taking as given the reactions of future policymakers.

This 'discretionary' method has been usefully interpreted by Oudiz and Sachs (1985), Currie and Levine (1985) and numerous subsequent authors as a Stackelberg leadership game, with the policymaker at each point in time treated as a separate actor, and playing the role of leader to all incarnations of its future self. In general the solution to this game will not see W_0 obtain its maximum value on the constraint set, for the same reason that the outcome of any non-trivial noncooperative game will not generally be the best conceivable for any one player: agents have different preferences over dynamic outcomes from any given time period onwards, so the policymaker at time t will not be considering exactly the same trade-off for that period as would the policymaker at time 0. A further complication is non-uniqueness: even in simple linear-quadratic models, if lagged endogenous variables feature in the constraint set (i.e., if Z has at least one dimension) then there is no guarantee that a *unique* Markov-stationary solution to the Stackelberg leadership game exists (that is, a solution that sets x as a time-invariant function of the state of the world, say $x(z, \varepsilon)$). This point is made in a general setting by Blake and Kirsanova (2008). The problem is compounded by the fact that non-Markov 'reputational' equilibria may additionally exist, under which private-sector plans may react to the current choices of the policymaker in accordance with 'trigger strategies' – as studied by Ireland (1997). In sum, discretionary outcomes are not particularly desirable for any policymaker, and may be extremely unpredictable.

The second method that has been commonly applied to deal with time inconsistency is to assume some form of institution exists at time 0 capable of augmenting the objective of the policymaker in all subsequent time periods, so as to incorporate a direct concern for the impact of current variables on past expectations. In similar problems to those presented in this paper, the 'Recursive Contracts' work of Marcet and Marimon (1998) has shown how augmentations to the objective function of the policymaker in each period can be used to incorporate the cost associated with violating past promises, where that cost is assessed from the perspective of the period in which those promises were originally made. To achieve this, future policymakers are contractually obliged to maximise the welfare function V_t rather than W_t in all time periods t, where $V_0 = W_0$ and V_t is given for $t > 0$ by:

$$V_t = W_t + \sum_{s=1}^{\min\{t,T\}} \beta^{-s} \lambda_g \left(h \backslash s, t-s\right)' g_s \left(x\left(h,t\right), \varepsilon\right), \tag{2.6}$$

where h is the history observed to time t, and the lagged shadow values $\lambda_g \left(h \backslash s, t-s\right)$ are set equal to the values they were assigned when optimising in period $t-s$. Straightforward comparison with optimality condition (2.4) reveals that repeated implementation of this objective will ensure we maximise W_0 subject to all relevant constraints. Marcet and Marimon show (under slightly more restrictive conditions on the forward-looking constraints[53]) that a single new state variable can be defined to store the relevant values of Lagrange multipliers as time progresses, ensuring the problem then has a recursive representation that permits dynamic programming techniques to be applied. When the forward-looking constraints are lifetime utility promises in insurance or moral hazard problems these multipliers have natural interpretations as equivalent to Pareto weights, applied in the recursive solution of a period-by-period weighted-utilitarian social planning problem – a point exploited in particular by Mele (2011).

[53] Specifically, either that $T = 1$ or that $T = \infty$. In the latter case they also require that $g_{s+r}\left(\cdot\right) = \beta^r g_s\left(\cdot\right)$ for all $r > 0$ and some $\beta \in (0,1)$, as in the case of lifetime utility constraints. These restrictions ensure a simple mapping exists from the multipliers attached to particular g functions from one period to the next.

But there is an asymmetry in this setup that a number of authors have found problematic: why should the policymaker at time 0 have access to institutions that permit him or her to impose on all future policymakers the objective function W_0, when these institutions are implicitly not available at future dates? Returning to the example of a deterministic 'inflation bias' problem, for instance, why is it that the policymaker who announces an initial, beneficial inflationary expansion – to be followed by future retrenchment – is uniquely able to break free from all prior commitments? Or, to quote Svensson (1999): "Why is period 0 special?"

In his seminal work on monetary policy, Woodford (2003) has advocated that policymakers overcome this complication by devolving choice to institutions that adopt what he dubs a 'timeless perspective' on optimality. This borrows from the Recursive Contracts literature the idea that the objective function should be augmented to incorporate the impact of current decisions on past constraints, but renders that augmentation itself time-invariant by assuming the welfare function that all choices are implicitly maximising is one that belonged to a policymaker in the *infinitely* distant past. That is, it imposes the objective:

$$V_t = W_t + \sum_{s=1}^{T} \beta^{-s} \lambda_g \left(h \backslash s, t-s\right)' g_s \left(x\left(h,t\right), \varepsilon\right), \tag{2.7}$$

for all $t \geq 0$ and histories h.

There remains here the problem of choosing values of $\lambda_g \left(h \backslash s, t-s\right)$ for $s > t$, which cannot be taken from any previous optimisation problem. Woodford advocates doing so by first solving the policy problem for generic values of these inherited multipliers, then setting the multipliers in accordance with any function of the history of exogenous and/or endogenous variables that is also satisfied by *subsequent* costate variables – that is, those determined endogenously within the optimisation problem. So we will have $\lambda_g \left(h,t\right) = \lambda_g \left(h,s\right) = \lambda_g \left(h\right)$ for all t, s, and h.

A slight problem with this prescription is that there are in general many admissible functions of this form. This is because lagged endogenous variables are themselves mappings from history along the dynamic path that follows from optimisation. For instance, in a canonical linearised New

Keynesian model the optimal choices for output and inflation will exhibit collinearity with one another. Dependence of optimal policy on lagged output and shocks can always be equivalently expressed as dependence upon lagged inflation and shocks, or some convex combination. Prior to the optimal policy regime being established, though, this collinearity need not hold – so two distinct 'timeless' policies will not (initially) be equivalent to one another. Particular criteria do exist that will deliver uniqueness – for instance, that optimal policy can subsequently be represented by a 'targeting rule' that is robust to misspecification of the underlying shock process (see Giannoni and Woodford (2002) for a full presentation), or that dependence upon lagged endogenous variables is limited to the 'state variables' in Z realised in the time period immmediately preceeding the present. But in general it is not clear which of these is the most appropriate to apply.

The 'timeless' approach also faces challenging problems from a normative perspective. Recalling once more that our basic problem was a lack of consistency between the objectives of policymakers at different horizons, it overcomes this issue by requiring that all policymakers adopt the hypothetical objective of a decisionmaker who never actually existed. As this interpretation suggests, one can construct examples in which a policy determined in such a fashion violates a simple Pareto criterion: one can identify two alternative paths for the model's endogenous variables, say A and B, such that the policymaker in every time period at least weakly prefers A to B (under objective W_t), and at least one policymaker *strictly* prefers A to B, but the timeless perspective nonetheless selects B over A. We provide an example in which this is true in Section 3.

If one is to permit meaningful choice for policymakers in each period whilst at the same time improving upon discretionary outcomes, the choice problems of those in at least *some* time periods must certainly be adjusted somehow; but it is very unsatisfactory to be doing this in a manner that does not guarantee that the chosen policy lies on the Pareto frontier. Notice by way of contrast that the recursive contract that enforces on future policymakers the objective W_0 *does* abide by the Pareto principle, since, by design, it maximises the welfare criterion of the policymaker at time 0.

One final important point regarding the timeless perspective approach is that it is not clear the objective (2.7) is well-defined in the event that $T = \infty$, even though the sum in (2.2) may well be. To ensure convergence one would need $\lim_{s\to\infty}\left(\frac{g_s\left(x\left(h^t,t\right),\varepsilon_t\right)}{\beta^s}\right) = 0$, but it is unlikely that this will hold in models of interest that *do* feature infinite sums of future variables in expectational constraints, since these sums very often relate to the net present value of revenue or utility streams, discounted at rate β (or, in the case of revenues, a gross interest rate whose inverse approaches β in steady state). An obvious example would be a fiscal policy problem in the style of Lucas and Stokey (1983), in which taxation to meet government expenditure needs must be spread optimally through time, implying optimisation subject to a net present value constraint on future primary surpluses. This boundedness problem does not arise under the Recursive Contracts objective (2.6), since there one only ever need consider the impact of current variables on constraints dating back to time 0; but even then a need to make good on past obligations could come to dominate policymaking as t becomes large.

2.3 Policymaking behind a veil of ignorance

This paper's focus is on policy problems for which relevant time-0 contracts cannot plausibly be written (or enforced). One example might be dynamic tax problems, in which future benefits may be promised as part of a current incentive structure: a democratic government cannot assume automatically that a chosen tax regime will endure in perpetuity, since no administration has the capacity to *enforce* this desire on its successors. Seeking the best 'time-0' policy under a rigid assumption that future generations will accept anything they may inherit does not then seem right – particularly if the implied long-run outcomes are notably *un*desirable for these future generations (as when 'immiseration' characterises limiting outcomes to dynamic income tax problems).

But we do not wish to rule out *some* kind of policy commitment. In practice central banks *are* given independence, tasked with pursuing specific, lasting objectives, and tax legislation *is* drafted with the intention that it may endure for many years. The reason for this presumably derives – at least in part – from the gains that policymakers at all horizons can attain by moving away from an uncooperative 'discretionary' equilibrium. The question is how this can be done in a manner that will at least be perceived 'just' by future generations – not treating their preferences as any less significant than those of their predecessors.

The most famous modern treatment of justice is due to Rawls (1971), and of particular applicability in our setting is his idea that lasting social institutions should be designed without reference to the *particular* economic (and social) circumstances of those whose task it is to construct them. According to the famous thought experiment, it should be as if a 'veil of ignorance' prevents individuals from knowing the specific place they are to take in the society whose institutions they are designing. In the class of models that we study, the only relevant 'particularity' of each policymaker is his or her position in time, relative to the period in which a given set of choice variables will be observed. For this reason, we suppose that a 'just' intergenerational settlement can be reached by choosing policy as if a veil of ignorance conceals knowledge of the time period in which it will apply. Policy is still allowed to be specific to prevailing economic *circumstances*, just not the time period *per se*.

In this regard our focus is on changing the information set used to determine policy, rather than changing underlying preferences. Notably, we do *not* obtain our results by changing the time preference parameter used to weight outcomes in distinct time periods. Applying a higher 'social discount factor' is an approach that has been taken by a number of authors sceptical about the possibility of implementing allocations that are optimal from the perspective of initial generations only. Examples include Phelan (2006) and Farhi and Werning (2007), both of whom focus on dynamic problems in which the optimal allocation from the perspective of initial generations results ultimately in immiseration for their descendents. Indeed, one of the main contributions of this paper is to separate the question of how to deal with time-inconsistency problems from the question of the appropriate social discount factor. Of particular interest in this context, we show in Section 5 that immiseration results cease to hold when policy is set without knowledge of time, even allowing the social discount factor to equal the private time preference parameter.

But a requirement that policy should not depend on time is very difficult to incorporate into the objective function directly. In particular, we need to do more than just deny the policymaker the capacity to condition on the time period when choosing the variables in x (though this is certainly necessary). This is because many other pieces of data could, in principle, still be used to *infer* the time period with great precision. For instance, if the policymaker is allowed to condition upon h_ε (the history of exogenous shocks) and still has knowledge of the period-0 realisation of this vector (as part of the history h^0, which is known under the original statement of the policy objective in (2.1)), then it will be a very simple procedure to match up the two histories and infer the number of extra periods included in h_ε relative to the initial shock history. A number of further steps are required to 'cleanse' the policy problem fully of time-specific information, and much of this section is devoted to presenting these steps in turn.

2.4 Equivalence classes defined

It is first useful to define an 'equivalence class' of the variables upon which x might depend, $H_e^\infty \times T_e \subset H^\infty \times \mathbb{Z}$, as a subset of the possible histories and dates such that under the chosen policy strategy we will have, for any two history-date pairs (h', t') and (h'', t''):

$$\left(h', t'\right), \left(h'', t''\right) \in H_e^\infty \times T_e \Longrightarrow x\left(h', t'\right) = x\left(h'', t''\right) \tag{2.8}$$

So if two history-date pairs 'share an equivalence class', this means that the same choice of endogenous variables must be made in response to either. An equivalence class is, therefore, a subset of the space $H^\infty \times \mathbb{Z}$ within which we assume the policymaker behind a veil of ignorance is unable to distinguish between distinct elements – and must set a common policy for all. In what follows it is also necessary to specify equivalence classes for choices of variables in Y or Z in isolation. Two history-date pairs

(h', t') and (h'', t'') 'share an equivalence class for control variables' if the following is true:

$$\left(h', t'\right), \left(h'', t''\right) \in H_e^{\infty} \times T_e \Longrightarrow y\left(h', t'\right) = y\left(h'', t''\right) \tag{2.9}$$

with a symmetric definition for state variables (elements of Z).

The purpose of these equivalence classes is to specify the precise subset of variables in $(h, t) \in H^{\infty} \times \mathbb{Z}$ about which we wish a policymaker choosing elements of Y and Z to be informed. If, for instance, two identical histories observed at different time periods always share an equivalence class, this is the same as saying that policy cannot be set as a function of t. This is certainly something we wish to impose, at least for the choice of control variables – to which the time-inconsistency problem attaches. The next sub-section states the requirement formally.

2.5 Time invariance

A simple time-invariance restriction on optimal choice over the variables in Y can be written as follows:

Condition 1 *For any $h \in H^{\infty}$ and all $s, t \geq 0$, (h, s) and (h, t) share an equivalence class for control variables.*

That is, if the history observed prior to period t were the same as that observed prior to period s, we would expect the same policy choice over Y in each of these periods. This is clearly a necessary restriction if choice over these variables is to be made in a way that is not particular to the preferences of a given period.[54]

But even in models for which Z is empty, this restriction is not enough by itself to answer the question: 'What would be done by a policymaker ignorant of time?' The information set remains far too rich, containing initial information about the history of the economy that could easily be used to

[54] It is not immediately clear that the same condition need apply for the choice of variables in Z, since the time-inconsistency problem in our setup derives entirely from expectational constraints, in which these variables do not feature.

infer the time period with near-perfect accuracy, even absent strict knowledge of t. For example, it would be a 'time-invariant' strategy in our framework to specify a distinct macroeconomic policy to be applied 23 years after the exogenous shock 'Germany is reunified' from that applied 24 years after the same shock. The constraint that the strategies in question would have to be replicated were Germany to be re-reunified is scarcely a restrictive one (history does not repeat itself *that* closely). What we want is a policymaker who acts as if understanding history's *events*, without recalling its *dates*; we have yet to ensure this.

To eliminate all possibility of time-specific choice in the general problem requires a number of complicated steps. It helps to develop the key intuition in as simple an environment as possible first. For this reason we restrict our attention temporarily to purely 'forward-looking' models – that is, models in which no constraints of the form (2.3) feature, and so $X = Y$. The following sub-section explores this simpler case, with the arguments subsequently generalised to models *with* state variables.

2.6 Purely forward-looking models

Our aim is to ensure state-contingent choices will be made without explicit or implicit knowledge of time. It is implicit knowledge (or inference) that is the harder to rule out. Even when time cannot be used directly as an argument in the policy function it may be possible to use other variables as fairly precise 'proxies'. To prevent this we must introduce two additional restrictions on the manner in which policy rules can be chosen: first, a restriction on the set of variables to which policy is permitted to be linked, which prevents choices from being used as vehicles to 'signal' time; second, a specification of the relative weightings that the policymaker must apply across different initial states of the world, chosen so that the likelihood of observing a given history will not depend upon time. These will also be the two key steps in ensuring time-independent choice when we reintroduce 'backward-looking' constraints, though the procedure there must be slightly more involved.

2.6.1 Restricting history dependence

The application of optimal policy will induce a singularity in the set of realised outcomes through time, in the sense that the successor history to any given $h \in H^{\infty}$ must feature as its last entry $x(h)$ (along with some drawing from the set of exogenous variables E). Moreover, the exact nature of this singularity is endogenous to the policymaker's choice, and so would be easy to exploit to infer time from history: given some information about the initial history of outcomes (a distribution across h^0), choice could be used to make sure certain realisations of $h \in H^{\infty}$ were (almost surely) particular to a known time period.[55] But note that this singularity arises from the presence of lagged endogenous variables only: the evolution of $h_{\varepsilon} \in E^{\infty}$ will, by definition, be unaffected by choice. Moreover, in a purely forward-looking model the variables in $h_x \in X^{\infty}$ are irrelevant to the basic structure of the problem going forward from every horizon: past values of lagged endogenous variables do not feature in the objective function, nor in the model's constraints.[56] Any desire to respond optimally to past shocks can always be incorporated directly through dependence on the shocks themselves, rather than through dependence on lagged endogenous variables that were in turn chosen in response those shocks. Thus the simplest way to prevent the policymaker from exploiting the singularity that choice imposes on history is to rule out dependence on lagged endogenous variables *per se*. We therefore apply the following:

Condition 2 *Suppose* $X = Y$. *For any* $t \geq 0$ *and all* $(h'_{\varepsilon}, h'_{y}), (h''_{\varepsilon}, h''_{y}) \in H^{\infty}$, *if* $h'_{\varepsilon} = h''_{\varepsilon}$ *then* $((h'_{\varepsilon}, h'_{y}), t)$ *and* $((h''_{\varepsilon}, h''_{y}), t)$ *share an equivalence class.*

In words, this condition says that for purely forward-looking models, if any two histories differ only in the realised values of *endogenous* variables at some point in the past then the same policy must be applied in response to both. If we apply Conditions 1 and 2 together in purely forward-looking models,

[55] As an extreme example, it would always be possible to signal the passage of time through the value assigned at an arbitrary decimal place for any one of the chosen variables.

[56] More generally this is always true of the variables in $h_y \in Y^{\infty}$, which is the same object in the forward-looking case.

endogenous variables must be chosen as functions of the history of exogenous variables alone, and we can denote these choices by $y(h_\varepsilon)$ (or $x(h_\varepsilon)$).

2.6.2 Weighting initial histories

Restricting dependence to exogenous variables prevents choice from exploiting policy-induced singularities, but it does not address the additional need to deny the policymaker knowledge of the initial history h^0, since h_ε^0 would still generally be enough information to tailor history-contingent policy actions to the time-periods in which those histories were more likely to be observed. We want the policymaker to act as if drawings from the space E^∞ are completely orthogonal to the realised time period: any given history is no more likely to characterise one period than another. We achieve this by assuming an information set that weights initial histories according to the *unconditional* distribution over E^∞ associated with the exogenous shock process, which we have called F. This unconditional distribution is characterised by the property:

$$\int_{E^\infty} \left[\int_{E^\infty} \eta(h_\varepsilon)\, dF_t\left(h_\varepsilon | h_\varepsilon^0\right) \right] dF\left(h_\varepsilon^0\right) = \int_{E^\infty} \eta(h_\varepsilon)\, dF(h_\varepsilon) \tag{2.10}$$

for all functions $\eta : E^\infty \to \mathbb{R}$. Only when initial histories are weighted according to this function is the likelihood of observing a given $h_\varepsilon \in E^\infty$ truly time-invariant for the policymaker. For this reason it is uniquely consistent with our desire to render the information set orthogonal to the time period.

Policy is thus chosen to maximise the objective W_0^y:

$$W_0^y = \int_{E^\infty} \left[\sum_{t=0}^{\infty} \beta^t \int_{E^\infty} \pi(x(h_\varepsilon), \varepsilon)\, dF_t\left(h_\varepsilon | h_\varepsilon^0\right) \right] dF\left(h_\varepsilon^0\right) \tag{2.11}$$

subject to constraint (2.2) holding for all realisations of h_ε. The superscript y is used here to denote the fact that this policymaker is choosing variables from the set Y only. For purely forward-looking models this is a trivial consequence of the fact that there are no variables in Z. For all cases of interest the integrals and sum contained in (2.11) will be well-behaved, and so we will have:

$$\begin{aligned} W_0^y &= \sum_{t=0}^{\infty} \beta^t \int_{E^\infty} \pi\left(x\left(h_\varepsilon\right), \varepsilon\right) dF\left(h_\varepsilon\right) \\ &= \frac{1}{1-\beta} \int_{E^\infty} \pi\left(x\left(h_\varepsilon\right), \varepsilon\right) dF\left(h_\varepsilon\right) \end{aligned} \tag{2.12}$$

The discount factor is irrelevant to this criterion: we could equivalently multiply by $(1-\beta)$ and preserve preference orderings. This is a consequence of two factors: the absence of any lagged endogenous variables in the constraints of the problem, and our explicit desire to set x according to a time-invariant function. Since there are no 'structural' linkages between outcomes in one time period and those of the preceeding period the dynamic specification of preferences ceases to be important; maximising the expected value of the criterion function π is the best strategy for *all* periods.

2.6.3 Characterising optimal policy

Definition 1 *Suppose $X = Y$. Then we will say that a policy function $x(h_\varepsilon)$ is 'optimal from behind a veil of ignorance' in the forward-looking case if it maximises the criterion W_0^y on the set of functions mapping from E^∞ to X such that the constraints (2.2) hold for all realisations of h_ε.*

This definition is reassuringly straightforward, which is precisely why we have chosen to focus initially on the forward-looking case. The general problem requires a number of extra considerations, resulting in a much more involved definition of optimality – but with the same basic principles at its core.

We now proceed to characterise optimal policy from behind a veil of ignorance, assuming for simplicity that the objective and constraint functions are differentiable. We use Lagrange multipliers, allowing for comparison with our earlier discussions of the time-0-optimal policy, and denote the present-value Lagrange multiplier on (2.2) at time t following history h_ε by $\lambda_g(h_\varepsilon)$. Making use of the iid assumption on the exogenous variables, first-order conditions to the problem may be written as:

$$\sum_{t=0}^{\infty}\left\{\beta^{t}\frac{\partial\pi\left(x\left(h_{\varepsilon}\right),\varepsilon\right)}{\partial y_{i}\left(h_{\varepsilon}\right)}+\sum_{r=0}^{\min\{t,T\}}\beta^{t-r}\lambda_{g}\left(h_{\varepsilon}\backslash r\right)'\frac{\partial g_{r}\left(x\left(h_{\varepsilon}\right),\varepsilon\right)}{\partial y_{i}\left(h_{\varepsilon}\right)}\right\} \tag{2.13}$$

$$=0$$

This condition must hold for F-almost all $h_{\varepsilon}\in E^{\infty}$, for all $y_{i}\left(h_{\varepsilon}\right)$ functions. Importantly, the absence of time-specific policy functions or distributions (in contrast with the earlier expression (2.4)) allows this expression to be rewritten to collect terms with common coefficients β^{t}:

$$\sum_{t=0}^{\infty}\beta^{t}\left\{\frac{\partial\pi\left(x\left(h_{\varepsilon}\right),\varepsilon\right)}{\partial y_{i}\left(h_{\varepsilon}\right)}+\sum_{r=0}^{T}\lambda_{g}\left(h_{\varepsilon}\backslash r\right)'\frac{\partial g_{r}\left(x\left(h_{\varepsilon}\right),\varepsilon\right)}{\partial y_{i}\left(h_{\varepsilon}\right)}\right\}=0 \tag{2.14}$$

This immediately implies:

$$\frac{\partial\pi\left(x\left(h_{\varepsilon}\right),\varepsilon\right)}{\partial y_{i}\left(h_{\varepsilon}\right)}+\sum_{r=0}^{T}\lambda_{g}\left(h_{\varepsilon}\backslash r\right)'\frac{\partial g_{r}\left(x\left(h_{\varepsilon}\right),\varepsilon\right)}{\partial y_{i}\left(h_{\varepsilon}\right)}=0 \tag{2.15}$$

2.6.4 Discussion

It is useful to contrast this result with that obtained by applying the 'timeless perspective' method. The first-order conditions to the problem of maximising the timeless perspective objective (2.7) subject to (2.2) can be expressed in our notation as:

$$\frac{\partial\pi\left(x\left(h,t\right),\varepsilon\right)}{\partial y_{i}\left(h,t\right)}+\sum_{r=0}^{T}\beta^{-r}\lambda_{g}\left(h\backslash r,t-r\right)'\frac{\partial g_{r}\left(x\left(h,t\right),\varepsilon\right)}{\partial y_{i}\left(h,t\right)}=0 \tag{2.16}$$

The main difference between expressions (2.15) and (2.16) is the presence of the inverse discount factors in (2.16).[57] This can be explained straightforwardly. Consistent with the interpretation of it as

[57] A more subtle difference is that the timeless perspective rule permits information on lagged endogenous variables to contribute to the determination of the multipliers and current endogenous variables in (2.16). As already discussed, this

doing what is best from the perspective of a time period in the distant past, the timeless perspective assumes that the marginal effect of relaxing a forward-looking constraint ought to be valued according to the relative priorities of a policymaker for whom that constraint was binding – under the assumption that that policymaker applied the same timeless perspective approach. This means weighing the 'direct' marginal effects on the criterion function $\pi(\cdot)$ at time t of changing the variable $y_i(h,t)$ against the 'indirect' marginal effects of relaxing constraints *at time* $t-r$ for some $r \in \{0,...,T\}$, regardless of whether $t \geq r$ holds (that is, regardless of whether the constraints in question ever actually bound).

The 'veil of ignorance' method, by contrast, assumes that constraints prior to time 0 are irrelevant, but that *a given history* h_ε *could just as easily characterise period* $t+r$ *as it could period* t. In the event that it did characterise $t+r$, the marginal value of relaxing the period-t forward-looking constraint that binds following history $h_\varepsilon \backslash r$ would be an appropriate object of concern for the policymaker, and this marginal value is given the same time-preference weighting as the marginal value of the *direct* effect of changing $y_i(h_\varepsilon)$ on the period-t criterion function $\pi(\cdot)$, that obtains in the event that h_ε instead characterises time t. So from behind a veil of ignorance we compare against one another the direct marginal benefits of changing variables at time t with the indirect marginal benefits of changing expectations of outcomes realised *subsequent* to t – even though it is (almost) never the case in practice that the same history h_ε would actually characterise both period t and period $t+r$.

Neither the 'veil of ignorance' approach nor the timeless perspective method is an inherently 'correct' means to resolve the time inconsistency problem. Both involve changing the policymaker's choice problem in order to ensure a time-invariant optimal choice. But the former does at least have the

will generally admit a multiplicity of timeless policy representations, each associated with a different set of values being given to the multipliers determined prior to time 0. These alternatives then converge on the same outcome over time, provided the effects of initial multiplier choices fade. In purely forward-looking models there will always exist a unique policy that satisfies the definition of timeless perspective policy *and* is independent of endogenous variables realised prior to time zero, but this is not always the policy emphasised in the literature (see, in particular, Giannoni and Woodford (2002)).

One of the advantages of the veil of ignorance approach over the timeless perspective is that all of the multipliers $\lambda_g(h_\varepsilon \backslash r)$ used in the former method are set endogenously as shadow values: there is no need to attribute fictitious value to the relaxation of constraints binding prior to time 0, and no multiplicity problem arises.

benefit that it is obtained via a sequence of restrictions deliberately constructed to consider the implications of denying the policymaker knowledge of time. These restrictions are therefore designed to approximate the normative principle of disinterested institutional design that many find appealing in the social contractarian literature. By contrast, the normative appeal of the principle: "Do what would have been in the best interests of our great-great-great-grandparents." is unclear at best. This is reinforced by the fact that in a very simple class of models – namely, the class of purely *deterministic* forward-looking models – the timeless perspective is Pareto-dominated by the policy that is best from behind a veil of ignorance. That is, no policymaker in any time period would ever prefer the timeless perspective strategy to the veil of ignorance one in this class of models. This point is illustrated in more detail in Section 3.

Condition (2.15), together with the model's constraints, will in general be sufficient to characterise a solution (given our regularity assumptions – notably concavity in the $\pi(\cdot)$ and $g(\cdot)$ functions). In the purely forward-looking models currently under focus this solution will also be an 'unconditionally optimal' policy, as defined by Damjanovic et al. (2008) (following Taylor (1979)): a time-invariant policy strategy that would maximise the *ex-ante* expected value of the policymaker's within-period objective criterion, where this expected value is in general considered at a sufficiently distant horizon from time 0 that a 'stochastic steady state' has been reached.[58] As Damjanovic et al. show, the unconditional expected value of steady-state welfare can be maximised by implementing optimality conditions equivalent to those associated with the timeless perspective, but with $\beta = 1$ – that is, with no time preference discounting. This is consistent with an exclusive focus on long-term outcomes, after a steady state has been reached. Denying the policymaker knowledge of time has a very similar effect, as the contrast between (2.15) and (2.16) makes clear.

[58] In the methodology of Damjanovic et al. there is implicit indifference across all paths that result in the same stochasatic steady state being reached, regardless how long it takes to get there. As we explain, veil of ignorance policy involves no transition process in purely forward-looking models: the distribution of outcomes (across E^{∞}) at time 0 under F is identical to the distribution at any subsequent horizon. In this regard it offers an obvious criterion for *refining* unconditionally optimal policy to obtain uniqueness – albeit one relevant only to the purely forward-looking case.

So in this class of models it just happens that the best policy from behind a veil of ignorance is also a time-invariant policy that maximises the unconditional expectation of *steady-state* welfare. But unlike Damjanovic et al we have not imposed a restriction that only outcomes in steady state *should* matter to the policymaker: the equivalence follows because *when* the policymaker is denied all information about time, endogenous variables must end up taking a time-invariant unconditional distribution from the outset. If they did not, the objective W_0^y could not be being maximised. The underlying unconditional distribution across exogenous variables, F, is used to weight different possible histories in the objective function, so it is unsurprising that the (induced) unconditional expectation across values of the criterion function $\pi(\cdot)$ is being maximised by veil of ignorance policy.

In the event that constraints additionally exist in the form of equation (2.3), veil of ignorance optimality should not necessarily imply unconditional optimality. Intuitively, the initial values z_{-1} that constrain choices at time 0 in the general case are not affected by chosen policy functions, and hence cannot be asserted *from the outset* to equal whatever function of lagged shock values they may ultimately satisfy in stochastic steady state. The method that we pursue allows for these variables to be chosen in a way that respects the structure of time preference, whilst still denying the policymaker knowledge of time *per se* (where appropriate). That is, our general setup avoids an exclusive focus on long-run outcomes. Indeed, a notable problem with the objective criterion adopted by Damjanovic et al. – as argued by Woodford (2010) – is that in the class of purely *backward*-looking problems (for which time-inconsistency is not an issue, since the relative rankings of policymakers across future state-contingent paths does not change through time) it recommends a strategy that differs from the usual optimal plan with discounting.[59] Its exclusive focus on steady-state outcomes is equivalent to assuming (counterfactually) that lagged endogenous variables appearing in backward-looking constraints at time 0 respond endogenously to the policymaker's time-0 choice.

[59] For instance, in a canonical Ramsey optimal savings problem the method would recommend a traditional 'golden rule', with the net marginal product of capital driven all the way to zero.

2.7 The general framework

2.7.1 Dividing the choice problem

We turn now to analysing the general case of models that additionally feature 'backward-looking' constraints, of the form (2.3). Central to our results in the previous sub-section was an assertion that a policymaker acting from behind a veil of ignorance ought to consider history h_ε as equally likely to occur in some time period t as in some distinct period s. We restricted dependence here to the history of exogenous variables alone as a deliberate measure to prevent the policymaker from exploiting the singularity that is induced by policy choices over the *joint* history of exogenous and endogenous variables – that is, to prevent the policymaker from specifying choices in a way that could subsequently permit the time period to be inferred (with some precision). It was additionally quite straightforward to select a distribution over initial histories that would imply no inference about time could be drawn from seeing any given $h_\varepsilon \in E^\infty$: this was the unconditional distribution F, which we could take as given by the priors of the model.

In the general case we can no longer avoid *some* dependence of our chosen plan upon lagged endogenous variables – specifically the 'state' variables in the set Z. This is because different realisations for these variables imply different sets of feasible policy choices. This introduces two problems. First, there is no reason why state variables could not also be set in a way that provided information to future policymakers about time. The same form of singularity between past history and current choice will obtain, with the same potential for choice to exploit this and 'signal' the time period.[60] Second, if policy can be conditioned on knowledge of state variables we will presumably need to make certain distributional assumptions regarding initial histories for these (just as we did when averaging over h_ε^0 under the distribution F in the purely forward-looking case), to prevent inference about time being drawn from an observed $z \in Z$. But there is no obvious distribution to apply: the priors of the model

[60] Again, a simple example would be to use as a signalling device otherwise irrelevant decimal places when specifying the value of a policy instrument.

do not provide us with an equivalent to F for the history of state variables, since these variables are endogenous to choice.

Together these problems make simultaneous choice over Y and Z infeasible. If control variables must now be set as a function of lagged states as well as the history of exogenous variables, it would always be possible to render certain (h_ε, z) pairs more likely to be observed in earlier than later time periods, no matter what distribution over initial histories is assumed. Similarly, choice over Y cannot be made *assuming* a fixed, stationary distribution over Z if the process that actually determines the evolution of the state variables is chosen simultaneously. For these reasons our strategy is to divide up the choice process, with separate (but simultaneous) problems solved to determine appropriate values of state and control variables. Specifically, we assume that a 'forward-looking' policymaker is given responsibility for choosing variables from the set Y in response to any given history, whilst a 'backward-looking' policymaker simultaneously chooses from the set Z. The precise informational restrictions imposed will differ between the two problems, reflecting our specific desire to preclude all knowledge of time *when selecting those variables to which a time-inconsistency problem attaches* – i.e., the variables in Y.[61] Conceiving of two distinct policymaking agents in this manner is intended as a presentational device rather than a delegation proposal per se: the idea is that the state-contingent policy obtained by dividing up the choice problem will be more appropriate for repeated implementation by an independent institution, or for long-term legislation, than a 'time-0 optimum' – because of its disinterested treatment of time.[62] The remainder of this sub-section develops the arguments in full.

[61] We ultimately show that policymakers at different points in time do not disagree in their preferences over variables in Z once those in Y are selected without knowledge of time. For this reason there is no disagreement to resolve, and so no necessity to deny knowledge of time when choosing in Z. Allowing knowledge of time when choosing in Z additionally makes the analysis more tractable.

[62] For example, the divided choice problem might be a way of answering the question: 'What is the appropriate inflation target?' This target itself could then be devolved to an independent central bank, without necessarily believing that that bank should represent the incarnation of either the forward- or the backward-looking policymaker that we study below.

2.7.2 Minimising dependence

Simply dividing choice over Y and Z between distinct policymakers does not itself rule out inter-linkages. Distinct policymakers can still signal to one another. The way that we will prevent this is directly analogous to our approach in the purely forward-looking case: we place limits on the amount of information available, and what it can convey. First, we will assume that history-contingent choice over the variables in Y must be made based on only a minimal knowledge of history. Second, we will ensure the policymaker choosing the variables in Y does so knowing only that some stationary distribution characterises the evolution of this history vector through time.

The first of these steps, minimising dependence, requires again that choice over the variables in Y is made without sufficient knowledge of the history of endogenous variables to be able to draw inference about time. This must certainly mean denying the 'forward-looking' policymaker information on his or her own past choices over Y. In general there is no loss to additionally requiring the forward-looking policymaker to condition only on the first lag of the variables in Z. Since the process generating state variables prior to the initial time period is completely unknown, and since only one lag of these variables actually constrains possibilities, allowing further dependence would introduce unnecessary arbitrariness into policy decisions.

For this reason we apply Condition 3 (a natural generalisation of Condition 2) in addition to the basic time-invariance restrition, Condition 1 (which, for obvious reasons, is retained throughout):

Condition 3 *For any $t \geq 0$ and all $\left(h'_\varepsilon, h'_y, h'_z\right), \left(h''_\varepsilon, h''_y, h''_z\right) \in H^\infty$, if $h'_\varepsilon = h''_\varepsilon$ and $z' = z''$ (where z' and z'' are the most recent entries in h'_z and h''_z respectively) then $\left(\left(h'_\varepsilon, h'_y, h'_z\right), t\right)$ and $\left(\left(h''_\varepsilon, h''_y, h''_z\right), t\right)$ share an equivalence class for control variables.*

The combined implementation of Conditions 1 and 3 will imply that the vector of control variables can be expressed as a function $y\left(h_\varepsilon, z\right)$, where z is the lagged state vector.

Limiting dependence is not just necessary for the control variables. When the choice problem is split up in the way we propose, a 'backward-looking' policymaker who knew that choices over Y would depend on past values in Z might have an incentive to exploit this, by deliberately *inducing* particular future realisations for the control variables through specific choices of the states – in a manner whose optimality need not be time-invariant. We do not want this to be a feature of choice: the whole point of dividing up the problem was that the control variables should be determined under the restricted information set of the 'forward-looking' policymaker alone. The easiest way to rule it out is for choice in Z to be made as a function of the shock history, the initial lagged state vector (i.e., z_{-1}), and the time period. These arguments are unaffected by the backward-looking policymaker's decisions, so to each point in the associated space $E^{\infty} \times Z \times \mathbb{Z}$ can be attached given values for the control variables without the backward-looking policymaker being able to influence the likelihood of distinct realisations within that space in response. We therefore impose the following Condition:

Condition 4 *For any $t \geq 0$ and all $\left(h'_{\varepsilon}, h'_{y}, h'_{z}\right), \left(h''_{\varepsilon}, h''_{y}, h''_{z}\right) \in H^{\infty}$, if $h'_{\varepsilon} = h''_{\varepsilon}$ and $z'_{-1} = z''_{-1}$ (where z'_{-1} and z''_{-1} are the entries in h'_{z} and h''_{z} respectively corresponding to time period -1) then $\left(\left(h'_{\varepsilon}, h'_{y}, h'_{z}\right), t\right)$ and $\left(\left(h''_{\varepsilon}, h''_{y}, h''_{z}\right), t\right)$ share an equivalence class for state variables.*

This condition implies the state variables will ultimately be chosen according to some function $z\left(h_{\varepsilon}, z_{-1}, t\right)$.

2.7.3 Linking the choice problems

To summarise, we are proposing that the variables in Y and Z should be determined by two quite distinct choice problems. The former (control variables) are to be set according to some function $y\left(h_{\varepsilon}, z\right)$, linking their values to the history of exogenous shocks and the lagged state vector. This is to be done under a supposition that state variables are evolving according to some stationary statistical model. The latter (state variables) are to be set according to a function $z\left(h_{\varepsilon}, z_{-1}, t\right)$, linking their values to the history of shocks, the initial lagged state vector, and the time period. This is to be done under a supposition that control variables are fixed in advance for each $\left(h_{\varepsilon}, z_{-1}, t\right)$ triple.

But a common set of constraints applies across the two problems, restricting the choices that can feasibly be made by any one of the agents, given the choices of the other. We need a way to incorporate these interactions into the choice setup appropriately. We do not want policy separation to extend so far that the forward-looking policymaker does not endogenise the effects on the backward-looking policymaker's problem of relaxing a particular constraint (and vice-versa). The sole purpose of separation is to allow information about time to be removed from all decisions affecting control variables, not to impose unnecessary dysfunctionality and coordination failures. For this reason we must adapt the Lagrangian method to our circumstances – allowing a common state-contingent shadow value to be attached to marginal relaxations of a given constraint, irrespective of which policymaker is choosing. The implication of this is that multipliers must feature in any representation of the pair of problems solved by, respectively, the forward- and backward-looking policymakers. A separate problem can then be defined to characterise these multipliers themselves.

2.7.4 Forward-looking policy: resolving an 'impossible trinity'

As in the purely forward-looking case, if the variables in Y are to be chosen absent any information about time we must additionally place certain restrictions on the forward-looking policymaker's initial distribution over histories – just as we applied the unconditional distribution F to weight realisations of h_{ε}^{0} in (2.11). This is more difficult when choice depends on lagged endogenous variables, since it is not clear exactly what distribution one should assume that these follow. Consistent with the discussion above, we do know that the policy function for control variables, $y : E^{\infty} \times Z \rightarrow Y$, will be chosen to solve a problem of the form:

$$\max_{y(h_\varepsilon,z)} \sum_{t=0}^{\infty} \beta^t \int_{E^\infty \times Z} \int_{E^\infty \times Z} \{\pi \left(x \left(h_\varepsilon, z\right), \varepsilon\right) \tag{2.17}$$
$$+ \lambda_g \left(h_\varepsilon, z\right)' \sum_{s=0}^{T} \int_{E^\infty \times Z} g_s \left(x \left(h'_\varepsilon, z'\right), \varepsilon'\right) dF_s \left(h'_\varepsilon, z' | h_\varepsilon, z\right)$$
$$+ \lambda_m \left(h_\varepsilon, z\right)' m \left(x \left(h_\varepsilon, z\right), z, \varepsilon\right)\} dF_t \left(h_\varepsilon, z | h^0_\varepsilon, z_{-1}\right) dF^0 \left(h^0_\varepsilon, z_{-1}\right)$$

for some $z : E^\infty \times Z \to Z$, $\lambda_g : E^\infty \times Z \to \mathbb{R}^N_+$ and $\lambda_m : E^\infty \times Z \to \mathbb{R}^M_+$ functions, for given horizon-specific conditional distributions $F_t \left(h'_\varepsilon, z' | h_\varepsilon, z\right)$ – which incorporate both the given conditional distribution over exogenous variables $F_t \left(h'_\varepsilon | h_\varepsilon\right)$ and *some* given conditional distribution over the variables in Z, denoted $F_t \left(z' | h'_\varepsilon, h_\varepsilon, z\right)$ – and for a given initial distribution $F^0 \left(h^0_\varepsilon, z_{-1}\right)$, which likewise incorporates both the unconditional distribution over shock histories, $F \left(h^0_\varepsilon\right)$, and *some* distribution over Z conditional on the initial history, $F^0 \left(z_{-1} | h^0_\varepsilon\right)$.[63]

The problem is in specifying two objects: the conditional distribution over (lagged) state variables at each horizon, $F_t \left(z' | h'_\varepsilon, h_\varepsilon, z\right)$, and the distribution over initial values of the lagged state vector, $F^0 \left(z_{-1} | h^0_\varepsilon\right)$. There are three considerations of relevance here:

1. **Stationarity:** We know that unless we choose the two distributions to satisfy stationarity (so that $\int F_t \left(h_\varepsilon, z | h^0_\varepsilon, z_{-1}\right) dF^0 \left(h^0_\varepsilon, z_{-1}\right) = F^0 \left(h_\varepsilon, z\right)$ holds) the forward-looking policymaker will be able to infer time with some degree of precision by knowing the $\left(h_\varepsilon, z\right)$ pair that obtains.

2. **Agreement with $z \left(h_\varepsilon, z\right)$:** It is natural to want the assumed function $z \left(h_\varepsilon, z\right)$ to 'agree' with the choice of conditional distribution, in the sense that $F_t \left(z' | h'_\varepsilon, h_\varepsilon, z\right)$ places positive probability mass only on those vectors in Z that could actually be reached by implementing the backward-

[63] We draw no distinction between the conditional distribution applied by the policymaker to assess the likelihood of future states, $F_t \left(h_\varepsilon, z | h^0_\varepsilon, z_{-1}\right)$, and the distribution $F_s \left(h'_\varepsilon, z' | h_\varepsilon, z\right)$ that is multiplying expectational constraint functions. This is equivalent to saying that the policymaker believes the private sector shares an identical stochastic model for the evolution of $\left(h_\varepsilon, z\right)$ to his or her own when forming its expectations, and applies shadow values accordingly.

looking policy function $z(h_\varepsilon, z)$ t times from an initial pair (h_ε, z) and with shocks evolving according to the relevant entries in h'_ε.

3. **Full support:** We would like the policymaker to place positive probability density on the possibility that *any* (h_ε, z) pair in the space $E^\infty \times Z$ could occur – otherwise choice in some regions of that space could be made completely arbitrarily, at no cost under the assumed policy objective.

In general these three objectives are not consistent with one another. To see this, suppose (plausibly) that an assumed $z(h_\varepsilon, z)$ function induces convergence to a stochastic steady state independent of initial conditions. That is, if we define a function $z^n : E^\infty \times Z \to Z$ such that $z^n(h_\varepsilon, z)$ is the mapping obtained by recursively applying the backward-looking policy function n times (starting from a pair $(h_\varepsilon \backslash (n-1), z)$ and allowing the shock history to evolve consistent with h_ε), suppose that there exists a function $z^*(h_\varepsilon)$ such that $\lim_{n\to\infty}[z^n(h_\varepsilon, z)] = z^*(h_\varepsilon)$. Then by the fact that we are converging on $z^*(h_\varepsilon)$, if $F_t(z'|h'_\varepsilon, h_\varepsilon, z)$ reflects the choices made under the $z(h_\varepsilon, z)$ function the only initial distribution consistent with stationarity is one that sets $P\left(z_{-1}|h^0_\varepsilon\right) = 1$ if $z_{-1} = z^*\left(h^0_\varepsilon\right)$ and $P\left(z_{-1}|h^0_\varepsilon\right) = 0$ otherwise. Any other distribution would have to be placing greater mass on values of z_{-1} different from $z^*\left(h^0_\varepsilon\right)$ than is placed on those values at more distant horizons following an identical history.[64] But this implies a singularity: the policymaker would be acting as if the stationary distribution had obtained from the very start of time, and thus would be placing no probability on the possibility of observing any pair $(h_\varepsilon, z) \in E^\infty \times Z$ such that $z \neq z^*(h_\varepsilon)$. This in turn would violate our third requirement: it would imply a completely arbitrary policy choice could be set for these regions in the parameter space, at no cost.

The basic problem is that if state variables are converging on a stochastic steady state it will be more likely for the forward-looking policymaker to observe values *away* from this steady state the earlier is the time period – unless we assume that these variables *start* in steady state, with certainty.

[64] That is, we could not have $\int F_t\left(z_{-1}|h^0_\varepsilon, h_\varepsilon, z\right) dF^0(z|h_\varepsilon) = F^0\left(z_{-1}|h^0_\varepsilon\right)$, since the object on the left-hand side of this expression must be converging to the degenerate distribution that places unit mass on $z^*\left(h^0_\varepsilon\right)$ as t becomes large.

But the latter would plainly not be an appropriate assumption in a great number of important models, where analysis of appropriate policy during the transition to steady state is important. We must violate at least one of the three requirements we had set for ourselves.

Since the focus of this Part is on the implications for policy of not knowing time, it is important to keep the stationarity condition if possible. If knowledge of the pair (h_ε, z) conveyed some information about time, there would remain an incentive to choose in a manner particular to the preference structure of the policymaker in period 0. It is also important to derive optimal strategies for the complete space $E^\infty \times Z$ if our method is to be useful for policy purposes. For this reason we do not wish to drop our third requirement, that the policymaker places positive probability mass on all proper subsets of $E^\infty \times Z$. The alternative would be arbitrariness away from steady state.

The second requirement – of consistency between the conditional distribution $F_t\left(z'|h'_\varepsilon, h_\varepsilon, z\right)$ and the backward-looking policy function $z\left(h_\varepsilon, z\right)$ – we choose to drop. This implies permitting a separation in the mind of the forward-looking policymaker between the stochastic process determining the evolution of *lagged* state variables through time, and the response of *contemporary* state variables to the state of the world through the function $z\left(h_\varepsilon, z\right)$. There is nothing inherently contradictory about doing this: mathematically the two objects $F_t\left(z'|h'_\varepsilon, h_\varepsilon, z\right)$ and $z\left(h_\varepsilon, z\right)$ are quite distinct. But there is still an awkwardness in the construct, which ultimately derives from the difficulty of the task we have set ourselves: how to choose optimal policy in response to variables that are converging on a stochastic steady state, without drawing inference from this convergence about the time period.

With stationarity satisfied we can carry out a simple manipulation to the forward-looking policymaker's objective, which becomes:

$$\max_{y\left(h_\varepsilon,x^b\right)} (1-\beta)^{-1} \int_{E^\infty \times Z} \{\pi\left(x\left(h_\varepsilon,z\right),\varepsilon\right) \\ +\sum_{s=0}^{T}\left[\int_Z \lambda_g\left(h_\varepsilon \backslash s, z'\right)' dF_{-s}\left(z'|h_\varepsilon \backslash s, h_\varepsilon, z\right)\right] g_s\left(x\left(h_\varepsilon,z\right),\varepsilon\right) \\ +\lambda_m\left(h_\varepsilon,z\right)' m\left(x\left(h_\varepsilon,z\right),z,\varepsilon\right)\} dF^0\left(h_\varepsilon,z\right) \tag{2.18}$$

where we define the conditional distribution $F_{-s}\left(z'|h_\varepsilon \backslash s, h_\varepsilon, z\right)$ analogously to $F_s\left(z'|h'_\varepsilon, h_\varepsilon, z\right)$, capturing the probability that the state vector took a value z' s periods ago, given a contemporary value for the lagged state vector of z and a current shock history h_ε.

What is particularly interesting about expression (2.18) is that the precise distribution function that we assume across (h_ε, z) pairs will only have an impact on optimal choice through its presence in the integral $\int_Z \lambda_g\left(h_\varepsilon \backslash s, z'\right)' dF_{-s}\left(z'|h_\varepsilon \backslash s, h_\varepsilon, z\right)$.[65] This integral gives the shadow values to the policymaker of relaxing s-period-ahead expectational constraints through changes in the control variables chosen when (h_ε, z) characterises the state of the world – assessed from the perspective of (and aggregated across) all prior situations for which the pair (h_ε, z) could conceivably characterise the state of the world s periods hence. The twin features (a) that only the *aggregate* of shadow values is needed to incorporate the value of relaxing expectational constraints, and (b) that the distribution function assumed by the forward-looking policymaker over state variables will be irrelevant except to the extent that it impacts upon this aggregate, are together very useful. They imply that the (arbitrary) distribution over Z will not matter at all for policy so long as we hold the *aggregate* of shadow values fixed for a given (h_ε, z) pair and given horizon s. As we have discussed, this distribution incorporates a stationarity assumption that is quite likely to be inconsistent with the assumed (and actual) backward-looking policy function $z\left(h_\varepsilon, z\right)$, so we do not wish it to play a significant role in the outcome of the policy process *per se*. The

[65] Note that the value of the main integral in the objective can be maximised piecewise, with the measure applied across it irrelevant for these purposes.

natural way to link the forward- and backward-looking problems is, then, to place cross-restrictions directly on this aggregate of shadow values – rather than making specific assumptions about both the distribution and about shadow values in specific states of the world. To save on notation we refer to the horizon-s aggregate relevant to policy choices in state of the world (h_ε, z) as $\lambda_g^s(h_\varepsilon, z)$:

$$\lambda_g^s(h_\varepsilon, z) = \int_Z \lambda_g\left(h_\varepsilon \backslash s, z'\right) dF_{-s}\left(z' | h_\varepsilon \backslash s, h_\varepsilon, z\right) \tag{2.19}$$

Hence the objective can be written:

$$\begin{aligned} &\max_{y\left(h_\varepsilon, x^b\right)} (1-\beta)^{-1} \int_{E^\infty \times Z} \{\pi\left(x\left(h_\varepsilon, z\right), \varepsilon\right) \\ &+ \sum_{s=0}^{T} \lambda_g^s\left(h_\varepsilon, z\right)' g_s\left(x\left(h_\varepsilon, z\right), \varepsilon\right) \\ &+ \lambda_m\left(h_\varepsilon, z\right)' m\left(x\left(h_\varepsilon, z\right), z, \varepsilon\right)\big\} dF^0\left(h_\varepsilon, z\right) \end{aligned} \tag{2.20}$$

Notice again the irrelevance of the discount factor to optimal choice within Y: this follows for identical reasons to those outlined in the purely forward-looking case.

2.7.5 Backward-looking policy

Consistent with the distinct equivalence class structure applied to it, the backward-looking policy problem is to choose a function $z : E^\infty \times Z \times \mathbb{Z} \to Z$ to solve:

$$\begin{aligned} &\max_{z\left(h_\varepsilon, z_{-1}, t\right)} \sum_{t=0}^{\infty} \beta^t \int_{E^\infty} \{\pi\left(x\left(h_\varepsilon, z_{-1}, t\right), \varepsilon\right) \\ &+ \lambda_g\left(h_\varepsilon, z_{-1}, t\right)' \sum_{s=0}^{T} \int_{E^\infty} g_s\left(x\left(h'_\varepsilon, z_{-1}, t+s\right), \varepsilon'\right) dF_s\left(h'_\varepsilon | h_\varepsilon\right) \\ &+ \lambda_m\left(h_\varepsilon, z_{-1}, t\right)' m\left(x\left(h_\varepsilon, z_{-1}, t\right), z\left(h_\varepsilon \backslash 1, z_{-1}, t-1\right), \varepsilon\right)\big\} dF_t\left(h_\varepsilon | h_\varepsilon^0\right) \end{aligned} \tag{2.21}$$

for given $y : E^{\infty} \times Z \times \mathbb{Z} \rightarrow Y$, $\lambda_g : E^{\infty} \times Z \times \mathbb{Z} \rightarrow \mathbb{R}_{+}^{N}$ and $\lambda_m : E^{\infty} \times Z \times \mathbb{Z} \rightarrow \mathbb{R}_{+}^{M}$ functions, the given distribution function over exogenous variables $F_t\left(h'_{\varepsilon}|h_{\varepsilon}\right)$, and any given set of initial conditions $\left(h_{\varepsilon}^{0}, z_{-1}\right)$.

The setup here is far more straightforward than for forward-looking choice, since the equivalence class structure now imposed implies that there is no need to specify a stochastic process for the evolution of state variables. These state variables are instead the objects of choice for each history-date pair (given the initial z_{-1} vector). The only stochastic process with which we are concerned relates to the exogenous variables, and is associated with a distribution function, F, that is known. Notice that the terms in the second summation are redundant for $s > 0$, since by definition they do not depend on any of the variables in Z. This implies that choice under objective (2.21) will always be time-consistent, since a given decision is relevant only to outcomes in contemporary and subsequent periods – not prior ones.

2.7.6 Determining the multipliers

The Lagrange multiplier functions remain to be determined in both of the above problems. We need to choose them in such a way that the constraints in the model will always be satisfied. We state the basic problem that these multipliers solve using their representation for the backward-looking policymaker's problem – that is, as functions of the triple $(h_{\varepsilon}, z_{-1}, t)$. Their values in the forward-looking problem will then be set by an appropriate mapping. The functions $\lambda_g : E^{\infty} \times Z \times \mathbb{Z} \rightarrow \mathbb{R}_{+}^{N}$ and $\lambda_m : E^{\infty} \times Z \times \mathbb{Z} \rightarrow \mathbb{R}_{+}^{M}$ are chosen, respectively, to solve the problems:

$$\min_{\lambda_g(h_{\varepsilon}, z_{-1}, t) \geq 0} \lambda_g\left(h_{\varepsilon}, z_{-1}, t\right)' \sum_{s=0}^{T} \int_{E^{\infty}} g_s\left(x\left(h'_{\varepsilon}, z_{-1}, t+s\right), \varepsilon'\right) dF_s\left(h'_{\varepsilon}|h_{\varepsilon}\right) \tag{2.22}$$

and

$$\min_{\lambda_m(h_{\varepsilon}, z_{-1}, t) \geq 0} \lambda_m\left(h_{\varepsilon}, z_{-1}, t\right)' m\left(x\left(h_{\varepsilon}, z_{-1}, t\right), z\left(h_{\varepsilon}\backslash 1, z_{-1}, t-1\right), \varepsilon\right) \tag{2.23}$$

for all $(h_\varepsilon, z_{-1}, t) \in H_\varepsilon^\infty \times Z \times \mathbb{Z}$, subject to given $z : E^\infty \times Z \times \mathbb{Z} \to Z$ and $y : E^\infty \times Z \times \mathbb{Z} \to Y$ functions and the exogenous conditional distribution function $F_s(h'_\varepsilon | h_\varepsilon)$.

This representation ensures, as usual, that there will be no (finite) solution for the multipliers so long as constraints are violated, and that multipliers will be set to zero whenever constraints are *strictly* satisfied.

2.7.7 Defining a solution

We are finally in a position to define formally optimal policy from behind a veil of ignorance in the general model. Because forward- and backward-looking variables are set through the choice of functions that are defined on distinct spaces from one another, the chief purpose this definition is to specify exactly how the two are linked. Before doing so it is useful to specify one preliminary definition: we will say that the conditional distribution $F_{-s}(z' | h_\varepsilon \backslash s, h_\varepsilon, z)$ 'agrees with' the policy function $z(h_\varepsilon, z)$ if it places strictly positive probability mass only on the $z' \in Z$ such that $z^s(h_\varepsilon, z') = z$.[66] The definition then takes two parts, relating respectively to the principal problems solved by the functions of interest, and the necessary links across the two evaluative spaces

Definition 2 *We consider the set of functions:* $y(h_\varepsilon, z_{-1}, t)$, $y(h_\varepsilon, z)$, $z(h_\varepsilon, z_{-1}, t)$, $z(h_\varepsilon, z)$, $\lambda_m(h_\varepsilon, z_{-1}, t)$, $\lambda_m(h_\varepsilon, z)$, $\lambda_g(h_\varepsilon, z_{-1}, t)$ *and* $\left\{\lambda_g^s(h_\varepsilon, z)\right\}_{s=0}^{T}$ *to characterise optimal policy from behind a veil of ignorance if and only if:*

1. *(a)* $y(h_\varepsilon, z)$ *solves (2.20) given* $z(h_\varepsilon, z)$, $\lambda_m(h_\varepsilon, z)$ *and* $\left\{\lambda_g^s(h_\varepsilon, z)\right\}_{s=0}^{T}$.

   (b) $z(h_\varepsilon, z_{-1}, t)$ *solves (2.21) given* $y(h_\varepsilon, z_{-1}, t)$, $\lambda_m(h_\varepsilon, z_{-1}, t)$ *and* $\lambda_g(h_\varepsilon, z_{-1}, t)$.

   (c) $\lambda_g(h_\varepsilon, z_{-1}, t)$ *solves (2.22) given* $y(h_\varepsilon, z_{-1}, t)$ *and* $z(h_\varepsilon, z_{-1}, t)$.

   (d) $\lambda_m(h_\varepsilon, z_{-1}, t)$ *solves (2.23) given* $y(h_\varepsilon, z_{-1}, t)$ *and* $z(h_\varepsilon, z_{-1}, t)$.

[66] Recall that $z^s : E^\infty \times Z \to Z$ is the mapping obtained by recursively applying the policy function $z(\cdot, \cdot)$ s times, starting from a pair $(h_\varepsilon \backslash (n-1), z)$ and with the shock history evolving according to the given $h_\varepsilon \in E^\infty$. We additionally define $z^0(h_\varepsilon, z) = z$ for convenience.

2. *The following equivalences hold for all* $(h_\varepsilon, z_{-1}, t) \in E^\infty \times Z \times \mathbb{Z}$ *and all* $s \in \{0, ..., \min\{t, T\}\}$.[67]

(a)

$$y(h_\varepsilon, z_{-1}, t) = y(h_\varepsilon, z(h_\varepsilon \backslash 1, z_{-1}, t-1), 0) \qquad (2.24)$$
$$= y(h_\varepsilon, z(h_\varepsilon \backslash 1, z_{-1}, t-1))$$

(b)

$$z(h_\varepsilon, z_{-1}, t) = z(h_\varepsilon, z(h_\varepsilon \backslash 1, z_{-1}, t-1), 0) \qquad (2.25)$$
$$= z(h_\varepsilon, z(h_\varepsilon \backslash 1, z_{-1}, t-1))$$

i.

$$\lambda_g(h_\varepsilon, z_{-1}, t) = \lambda_g(h_\varepsilon, z(h_\varepsilon \backslash 1, z_{-1}, t-1), 0) \qquad (2.26)$$

ii.

$$\lambda_g^s(h_\varepsilon, z) = \int_Z \lambda_g(h_\varepsilon \backslash s, z_{-1}, s)\, dF_{-s}^z(z_{-1} | h_\varepsilon \backslash s, h_\varepsilon, z) \qquad (2.27)$$

for some distribution function $F_{-s}^z(z_{-1} | h_\varepsilon \backslash s, h_\varepsilon, z)$ *that agrees with* $z(h_\varepsilon, z)$.

(c)

$$\lambda_m(h_\varepsilon, z_{-1}, t) = \lambda_m(h_\varepsilon, z(h_\varepsilon \backslash 1, z_{-1}, t-1), 0) \qquad (2.28)$$
$$= \lambda_m(h_\varepsilon, z(h_\varepsilon \backslash 1, z_{-1}, t-1))$$

Part 1 here defines the main problems that we wish the policy and multiplier functions to solve, whilst part 2 (in its multiple sub-parts) ensures consistency across the two different representations of these functions. Specifically, part 2 places two distinct requirements on each of the policy functions.

[67] We normalise $z(h_\varepsilon \backslash 1, z_{-1}, -1) \equiv z_{-1}$ where appropriate.

The first equality in each of its four parts (including part i of (c)) states that these functions should take a recursive form when defined on the triple $(h_\varepsilon, z_{-1}, t)$. That is, the policy implemented at horizon t following a shock history h_ε and an initial condition z_{-1} should be the same as the policy that *would have been* implemented in the initial time period following the same shock history h_ε and an initial condition equal to the state vector that was actually chosen at $t-1$. The second equality for parts (a), (b) and (d) then simply states that this form is carried over to the relevant policy function defined on the restricted space $E^\infty \times Z$.

Part 2 (c) ii requires more detailed discussion. It provides the crucial link between the multipliers on 'forward-looking' constraints across the two problems. Two points are important here. First, note that the equation characterises the *aggregate* values of the relevant multipliers, as used in objective (2.20). It does not specify independently shadow values that could be used in the original forward-looking problem (2.17) *for all prior states* for which an s-horizon constraint could be affected by changes to policy in the (subsequent) state (h_ε, z) – together with some distribution across these prior states. It states only that the aggregate of these shadow values – which can be inserted directly into (2.20) – must be equal to an equivalent aggregate obtained from the backward-looking problem. This relates back to our discussion of the 'impossible trinity' when determining appropriate forward-looking policy: It is quite possible that the backward-looking policy function $z(h_\varepsilon, z)$ induces convergence in the vector z on a stochastic steady state over time, where this steady state is described by some function $z^*(h_\varepsilon)$. If so, it is impossible for the forward-looking policymaker to assume a distribution over $E^\infty \times Z$ that is stationary, consistent with $z(h_\varepsilon, z)$, and has full support on $E^\infty \times Z$. Since we wish to retain stationarity and full support, we cannot hope for consistency between the *individual components* of the objects $\int_Z \lambda_g (h_\varepsilon \backslash s, z_{-1}, s)\, dF^z_{-s}(z_{-1}|h_\varepsilon \backslash s, h_\varepsilon, z)$ (which features in the preceeding definition) and $\int_Z \lambda_g (h_\varepsilon \backslash s, z')\, dF_{-s}(z'|h_\varepsilon \backslash s, h_\varepsilon, z)$ (which features in objective (2.18)), if the distribution $F^z_{-s}(z_{-1}|h_\varepsilon \backslash s, h_\varepsilon, z)$ is in turn consistent with a policy function that induces convergence. We can only hope to match the aggregate values of these objects.

The second aspect of part 2 (c) ii in the definition that requires some explanation is the distribution $F^z_{-s}(z_{-1}|h_\varepsilon\backslash s, h_\varepsilon, z)$ itself. This specifies the probability that any given $z_{-1} \in Z$ could have featured as the lagged state vector in the economy s periods ago, conditional upon the contemporary lagged state vector being z and the contemporary shock history being h_ε. We have placed a natural consistency requirement on this distribution – that it should put positive probability only on values of z_{-1} from which it is actually possible to reach z through s iterations of the backward-looking policy function.

We should stress that *in the vast majority of models this will be enough to characterise F^z fully.* Unless Z is discrete in certain dimensions it seems quite unlikely that two distinct 'starting values' for the state vector would result in precisely the same 'end value' obtaining s periods later (for a common shock history). But nothing in the setup that we have presented rules out the possibility of multiple starting values, so it is important to think a little about how F^z might be chosen in this event.

Applying Bayes' law to the probabilities associated with the measure F^z (where $P^z_s(z|h_\varepsilon, h_\varepsilon\backslash s, z')$ is used to denote the probability of observing state z s periods after observing state z', when history evolves according to h_ε) we have:[68]

$$P^z_{-s}\left(z'|h_\varepsilon\backslash s, h_\varepsilon, z\right) = \frac{P^z_s(z|h_\varepsilon, h_\varepsilon\backslash s, z')\,P^z(z'|h_\varepsilon\backslash s)}{\sum_{z''\in Z} P^z_s(z|h_\varepsilon, h_\varepsilon\backslash s, z'')\,P^z(z''|h_\varepsilon\backslash s)} \tag{2.29}$$

This involves an 'unconditional' measure, P^z,[69] detailing the prior probability attached to any given value of the state vector being observed as a 'starting value', s periods prior to the realisation of the z vector of interest.[70] If there is a unique z' such that $P^z_s(z|h_\varepsilon, h_\varepsilon\backslash s, z'') \neq 0$ (as when the mapping $z(h_\varepsilon, \cdot)$ is one-to-one) then the unconditional measure is irrelevant to P^z_{-s}, which takes a value of 1 for the unique z_{-1} consistent with evolving to z in s periods, and 0 elsewhere. But when multiple starting

[68] We work with probabilities here because of the likelihood that only a finite number of z' vectors will be consistent with evolution to a given z.

[69] That is, not conditional on prior or subsequent values of z.

[70] This is not quite the same as a prior distribution over *initial* (time -1) state vectors, since $P^z_{-s}(z'|h_\varepsilon\backslash s, h_\varepsilon, z)$ captures the probability of z' being a starting value for *any* s-period evolution to z when the shock history is h'_ε – starting from *any* time period from 0 onwards. Hence the distribution $P^z(z'|h_\varepsilon\backslash s)$ gives the probability of observing z' in any time period, given a shock history $h'_\varepsilon\backslash s$.

values can lead to the same end value for the state vector, the denominator of the fraction on the right-hand-side will generally exceed the numerator, and the exact specification of P^z will come to influence policy (it will determine how much weight is given to distinct multipliers from the backward-looking problem in the mapping (2.27)). The choice of this prior distribution would have to be made in a manner that respected our desire not to give specific emphasis to the initial conditions that happen to obtain at time 0 in any particular case (simply placing an atom on the *realised* value of z_{-1} would seem wrong), though beyond this it is hard to say much that is concrete. And since in this thesis we do not study an example for which a one-to-one mapping from past to future states does *not* apply, we are content to leave the matter open, pending a further understanding of cases in which it may prove important – acknowledging that our definition of veil of ignorance policy may need fine-tuning were such cases to arise.

2.7.8 Characterising a solution

We now proceed to derive optimality conditions associated with veil of ignorance policy. Given the regularity assumptions that we have made (in particular continuous differentiability), this can be done simply by taking piecewise first-order conditions to the distinct problems. We first consider the forward-looking policy problem: a necessary and sufficient condition for the function $y(h_\varepsilon, z)$ to solve this is that the following should hold for all y_i and all $(h_\varepsilon, z) \in E^\infty \times Z$:

$$\begin{aligned} &\frac{\partial \pi\left(x\left(h_\varepsilon, z\right), \varepsilon\right)}{\partial y_i\left(h_\varepsilon, z\right)} + \lambda_m\left(h_\varepsilon, z\right)' \frac{\partial m\left(x\left(h_\varepsilon, z\right), z, \varepsilon\right)}{\partial y_i\left(h_\varepsilon, z\right)} \\ &+ \sum_{s=0}^{T} \lambda_g^s\left(h_\varepsilon, z\right)' \frac{\partial g_s\left(x\left(h_\varepsilon, z\right), \varepsilon\right)}{\partial y_i\left(h_\varepsilon, z\right)} \\ &= 0 \end{aligned} \tag{2.30}$$

Consistent with our earlier discussion, note that this expression is independent of the discount factor β. Indeed, it is almost identical to condition (2.15) that we obtained in the purely forward-looking case, with the two exceptions (a) that the effects of the variables in Y on backward-looking constraints must now be considered, and (b) that the multipliers associated with the effects of changing policy in state (h_ε, z) on prior expectations can no longer be expressed simply by $\lambda_g(h_\varepsilon \backslash s)$, and instead must be collected in the more complex object $\lambda_g^s(h_\varepsilon, z)$.

Once more, the intuition for β not featuring is that changes to the value of $y(h_\varepsilon, z)$ are not considered by the forward-looking policymaker as being any more likely to impact upon one time period than another, so that if expectations s periods ahead matter to what can be achieved in time t, the effects of improving t-dated outcomes directly should be offset against the effects of relaxing t-dated constraints by changing outcomes at $t+s$. Even though it may be impossible for (h_ε, z) in reality to characterise *both* t and $t+s$, probabilistically it does (and with *equal* likelihood). Since the same considerations will apply symmetrically for all $t \geq 0$ (for any given t, all $s \in \{0, ..., T\}$ are relevant horizons – in contrast with the 'full commitment' consideration only of *past* constraints that may have bound), the discount factor reduces to an irrelevant constant in the objective function.

We next turn to the backward-looking policy problem: a necessary and sufficient condition for the function $z(h_\varepsilon, z_{-1}, t)$ to solve this is that the following should hold for all z_i and all $(h_\varepsilon, z_{-1}, t) \in E^\infty \times Z \times \mathbb{Z}$:

$$\begin{aligned}
&\frac{\partial \pi(x(h_\varepsilon, z_{-1}, t), \varepsilon)}{\partial z_i(h_\varepsilon, z_{-1}, t)} + \lambda_g(h_\varepsilon, z_{-1}, t)' \frac{\partial g_0(x(h_\varepsilon, z_{-1}, t), \varepsilon)}{\partial z_i(h_\varepsilon, z_{-1}, t)} \\
&+ \lambda_m(h_\varepsilon, z_{-1}, t)' \frac{\partial m(x(h_\varepsilon, z_{-1}, t), z(h_\varepsilon \backslash 1, z_{-1}, t-1), \varepsilon)}{\partial z_i(h_\varepsilon, z_{-1}, t)} \\
&+ \beta \int_{E^\infty} \lambda_m(h'_\varepsilon, z_{-1}, t+1)' \frac{\partial m(x(h'_\varepsilon, z_{-1}, t+1), z(h_\varepsilon, z_{-1}, t), \varepsilon')}{\partial z_i(h_\varepsilon, z_{-1}, t)} dF_1(h'_\varepsilon | h_\varepsilon) \\
&= 0
\end{aligned} \tag{2.31}$$

Notice that the condition depends on the time period t only to the extent that the functions $y(h_\varepsilon, z_{-1}, t)$, $\lambda_g(h_\varepsilon, z_{-1}, t)$ and $\lambda_m(h_\varepsilon, z_{-1}, t)$ do likewise: as usual, backward-looking policy is not subject to any time-inconsistency problem. Given the equivalences in part 2 of the definition of veil of ignorance optimality, an immediate corollary is that the following must hold for functions defined on the restricted space $E^{\infty} \times Z$:

$$\begin{aligned}
&\frac{\partial \pi\left(x\left(h_{\varepsilon}, z\right), \varepsilon\right)}{\partial z_i\left(h_{\varepsilon}, z\right)}+\lambda_g^0\left(h_{\varepsilon}, z\right)^{\prime} \frac{\partial g_0\left(x\left(h_{\varepsilon}, z\right), \varepsilon\right)}{\partial z_i\left(h_{\varepsilon}, z\right)} \\
&+\lambda_m\left(h_{\varepsilon}, z\right)^{\prime} \frac{\partial m\left(x\left(h_{\varepsilon}, z\right), z, \varepsilon\right)}{\partial z_i\left(h_{\varepsilon}, z\right)} \\
&+\beta \int_{E^{\infty}} \lambda_m\left(h_{\varepsilon}^{\prime}, z\left(h_{\varepsilon}, z\right)\right)^{\prime} \frac{\partial m\left(x\left(h_{\varepsilon}^{\prime}, z\left(h_{\varepsilon}, z\right)\right), z\left(h_{\varepsilon}, z\right), \varepsilon^{\prime}\right)}{\partial z_i\left(h_{\varepsilon}^{\prime}, x^{b \prime}\right)} d F_1\left(h_{\varepsilon}^{\prime} | h_{\varepsilon}\right) \\
&=0
\end{aligned} \tag{2.32}$$

Writing the backward-looking first-order condition in this form will allow us to neglect functions defined on the triple $(h_{\varepsilon}, z_{-1}, t)$ in subsequent analysis.

Optimal choices for the multiplier functions imply standard complementary slackness conditions, so that for all $(h_{\varepsilon}, z_{-1}, t) \in E^{\infty} \times Z \times \mathbb{Z}$ we have:

$$\lambda_g\left(h_{\varepsilon}, z_{-1}, t\right)^{\prime} \sum_{s=0}^{T} \int_{E^{\infty}} g_s\left(x\left(h_{\varepsilon}^{\prime}, z_{-1}, t+s\right), \varepsilon^{\prime}\right) d F_s\left(h_{\varepsilon}^{\prime} | h_{\varepsilon}\right)=0 \tag{2.33}$$

and

$$\lambda_m\left(h_{\varepsilon}, z_{-1}, t\right)^{\prime} m\left(x\left(h_{\varepsilon}, z_{-1}, t\right), z\left(h_{\varepsilon} \backslash 1, z_{-1}, t-1\right), \varepsilon\right)=0 \tag{2.34}$$

As usual, either the constraints must bind or the multipliers must equal zero. If we assume for simplicity that a one-to-one mapping from past to current states is always induced by the $z\left(h_{\varepsilon}, z\right)$ function (for all h_{ε}) then condition (2.33) may be expressed on the reduced space $E^{\infty} \times Z$ as:

$$\lambda_g^0\left(h_{\varepsilon}, z\right)^{\prime} \sum_{s=0}^{T} \int_{E^{\infty}} g_s\left(x\left(h_{\varepsilon}^{\prime}, z^{\prime}\right), \varepsilon^{\prime}\right) d F_s\left(h_{\varepsilon}^{\prime}, z^{\prime} | h_{\varepsilon}, z\right)=0 \tag{2.35}$$

with:

$$\lambda_g^s(h_\varepsilon, z) = \lambda_g^0(h_\varepsilon \backslash s, z') \quad if \quad P_{-s}^z(z'|h_\varepsilon \backslash s, h_\varepsilon, z) = 1 \tag{2.36}$$

for all $s \in \{0, ..., T\}$, or equivalently:

$$\lambda_g^s(h_\varepsilon, z) = \lambda_g^{s-1}(h_\varepsilon \backslash 1, z') \quad if \quad P_{-1}^z(z'|h_\varepsilon \backslash 1, h_\varepsilon, z) = 1 \tag{2.37}$$

for all $s \in \{1, ..., T\}$, where the distributions applied in the last three expressions agree with the policy function $z(h_\varepsilon, z)$.

When the policy mapping from past to current states is not one-to-one, (2.35) must still hold. This follows because we know the forward-looking constraints will still bind going forward from each (h_ε, z) pair in $E^\infty \times Z$ – since they certainly bind for each *triple* of the form $(h_\varepsilon, z, 0)$ when we work on the expanded space $E^\infty \times Z \times \mathbb{Z}$. The value of $\lambda_g^0(h_\varepsilon, z)$ must additionally equal $\lambda_g(h_\varepsilon, z, 0)$, by our definition of a solution. But the evolution of the multipliers no longer takes a simple recursive form. Instead we have the more general expression:

$$\lambda_g^s(h_\varepsilon, z) = \sum_{z' \in Z} \lambda_g^{s-1}\left(h_\varepsilon \backslash 1, z'\right) P_{-1}^z\left(z'|h_\varepsilon \backslash 1, h_\varepsilon, z\right) \tag{2.38}$$

for all $s \in \{1, ..., T\}$. As mentioned above, the choice of prior over starting values (z') is likely to matter for this expression. It may be simplest in practice to specify a prior that *does* associate a unique z' with each z (for a given h_ε), reverting back to condition (2.37). But it is debatable whether this would be possible whilst retaining our parsimonious perspective regarding the economic outcomes to which policy is intended to respond.

Finally, we can re-write the condition on the backward-looking multipliers for the reduced space $E^\infty \times Z$:

$$\lambda_m(h_\varepsilon, z)' m(x(h_\varepsilon, z), z, \varepsilon) = 0 \tag{2.39}$$

2.7.9 Analysis: asymmetric discounting

We can gain insight into the consequences of adopting our 'veil of ignorance' perspective if we compare its results with the first-order condition from the more traditional problem of maximising W_0 under full information about history and time. For the choice of generic variable x_i (which may be a component of either Y or Z), selected for history $h \in H^{\infty}$ and time period t, that first-order condition was given by equation (2.4), reproduced below:

$$\begin{aligned}&\frac{\partial \pi\left(x\left(h,t\right),\varepsilon\right)}{\partial x_i\left(h,t\right)}+\sum_{s=0}^{\min\{t,T\}}\beta^{-s}\lambda_g\left(h\backslash s,t-s\right)'\frac{\partial g_s\left(x\left(h,t\right),\varepsilon\right)}{\partial x_i\left(h,t\right)}\\&+\lambda_m\left(h,t\right)'\frac{\partial m\left(x\left(h,t\right),z\left(h\backslash 1,t-1\right),\varepsilon\right)}{\partial x_i\left(h,t\right)}\\&+\beta\int\lambda_m\left(h',t+1\right)'\frac{\partial m\left(x\left(h',t+1\right),z\left(h,t\right),\varepsilon'\right)}{\partial x_i\left(h,t\right)}dF_1\left(h'|h\right)\\&=0\end{aligned} \tag{2.40}$$

Consolidating conditions (2.30) and (2.32) into a single optimality requirement (holding for the choice of any $x_i\left(h_\varepsilon,z\right)$, whether in Y or Z), a sufficient condition for optimal policy from behind a veil of ignorance is that for every $\left(h_\varepsilon,z\right)\in E^{\infty}\times Z$ we have:

$$\begin{aligned}&\frac{\partial \pi\left(x\left(h_\varepsilon,z\right),\varepsilon\right)}{\partial x_i\left(h_\varepsilon,z\right)}+\sum_{s=0}^{T}\lambda_g^s\left(h_\varepsilon,z\right)'\frac{\partial g_s\left(x\left(h_\varepsilon,z\right),\varepsilon\right)}{\partial x_i\left(h_\varepsilon,z\right)}\\&+\lambda_m\left(h_\varepsilon,z\right)'\frac{\partial m\left(x\left(h_\varepsilon,z\right),z,\varepsilon\right)}{\partial x_i\left(h_\varepsilon,z\right)}\\&+\beta\int\lambda_m\left(h'_\varepsilon,z\left(h_\varepsilon,z\right)\right)'\frac{\partial m\left(x\left(h'_\varepsilon,z\left(h_\varepsilon,z\right)\right),z\left(h_\varepsilon,z\right),\varepsilon'\right)}{\partial x_i\left(h_\varepsilon,z\right)}dF_1\left(h'_\varepsilon|h_\varepsilon\right)\\&=0\end{aligned} \tag{2.41}$$

There are two notable differences between conditions (2.40) and (2.41). First, the form of (2.41) is time-invariant, whereas the number of terms in the summation in (2.40) depends upon t: that is, no time-inconsistency problem attaches to policy that is optimal from behind a veil of ignorance. When we are choosing (time-inconsistent) optimal policy from the perspective of time 0 to be implemented in a time period $t < T$, the impact of current choices on expectations formed prior to time 0 are irrelevant. But when we are choosing *absent knowledge of the time period for which that choice is being made* there are no choices to which some or all of the expectation channels can be deemed irrelvant: the possibility that a given (h_ε, z) pair will characterise any given $t \geq 0$ must be incorporated into each choice symmetrically.

Second, there is no 'inverse discounting' associated with expectational channels when policy is set from behind a veil of ignorance – that is, there are no β^{-s} coefficients in the summation in (2.41). Again, the logic here carries over from the forward-looking case. If we know the time period for which we are choosing, we can also associate a time period to any expectational consequences of that choice. So if, for instance, we have $s \leq T < t$ then we know that the (time-specific) choice of $x(h, t)$ affects policy possibilities at time $t - s$ through an expectation channel. Given the policymaker's relative time preferences, the marginal effects operating through this particular channel should be weighted by β^{-s}, relative to any 'direct' effects on time-t welfare (that is, relative to any marginal changes to the within-period criterion function π at time t) – since the associated benefits accrue s periods earlier. In contrast, when the policymaker is completely uninformed about the time period for which choice is being made there is an equal chance that any given history will characterise time t *or* time $t + s$ (even if there is zero probability that it will characterise both). Since the expectational effects operating at time $t + s$ have consequences for welfare at time t, one should give equal weight to the associated marginal effects. And since we do not care about effects on expectations formed prior to time 0, the first period (in our current example) for which the expectational consequences of choice are of relevance is period s. So we are concerned about direct effects in time periods $\{0, 1, 2, ...\}$ and expectational effects in periods $\{s, s + 1, s + 2, ...\}$, where the latter in turn feed through into *welfare* through relaxed constraints in periods $\{0, 1, 2, ...\}$. It is this symmetry that renders time preference weighting redundant in the choice of forward-looking variables.

But there is also an interesting *a*symmetry here. Optimal policy from behind a veil of ignorance does still imply that the welfare effects of changing variables in Z on policy possibilities one period hence *should* be discounted relative to any direct effects on the criterion function π.[71] This is consistent in particular with our focus on eliminating time-specific choice considerations as distinct from time preference discounting.

Any change to policy conceived at time 0 cannot possibly change the inherited state vector z_{-1}. The variables in Z can be affected by choices made in time periods $\{0, 1, 2, ...\}$, but with effects on backward-looking constraints that apply only in periods $\{1, 2, 3, ...\}$. So for each 'direct' effect on the criterion function π there is an effect via backward-looking constraints that should be weighted by the time preference parameter β relative to it. This is certainly a desirable feature of the approach: the specific time period is well known to be irrelevant to choice in a 'purely backward-looking' model, and since our general framework nests models without any expectational constraints at all there is no reason why veil of ignorance optimality should differ from standard solutions in these environments.

What the approach thus achieves is to find a middle way between the 'timeless perspective' method for finding a time-invariant policy rule, as outlined by Woodford (2003), and the method proposed by Taylor (1979) and, more recently, Damjanovic, Damjanovic and Nolan (2008) – of directly maximising the ('unconditional') expectation of *steady-state* welfare. As explained above, the former allows for choice to be time-specific, but supposes that it is designed to be optimal from the perspective of a policymaker who existed in the infinitely distant past. The problem of condition (2.40) being time-varying is this way overcome: the upper limit of the summation becomes T for all t. But the normative justification is shaky: why *should* we care about expectations formed prior to the new policy regime? Or in terms of our discussion above, considering the s-horizon expectational effect, why

[71] Even though we place few informational restrictions on the backward-looking policymaker, the choice of these variables remains time consistent – since the same set of constraints applies to them whether they are chosen in advance or contemporaneously.

should we associate each direct welfare gain at time t with an impact on expectations at time $t-s$? This implies a sequence of direct effects in periods $\{0, 1, 2, ...\}$ and expectational effects in periods $\{-s, -s+1, -s+2, ...\}$. The latter does not seem appropriate for $s > 0$. Moreover, from a technical perspective the timeless perspective encounters substantial difficulties when T is infinite, which occurs in many models of interest that incorporate 'lifetime' (dynastic) utility constraints. In that setting $\lim_{T\to\infty}\left(\beta^{-T}\frac{\partial g_T(x(h,t),\varepsilon)}{\partial x_i(h,t)}\right)$ is generally bounded away from zero, so that the summation in the timeless perspective variant upon (2.40) will not converge.[72]

The 'unconditional' optimisation approach, meanwhile, places all weight on steady-state outcomes. Since these are not necessarily reached in finite time, it implicitly allows either that there is no discounting for pure time preference, or that the policymaker counterfactually believes the distribution across initial values in Z is endogenous to their choice – so that any policy which ensures a particular $z \in Z$ is given greater weight under the *long-run* (induced) distribution function also renders that z more likely to characterise time period -1. The latter property is justifiably criticised by Woodford (2010): one can make a *normative* case against discounting for the sake of pure time preference (such as that famously delivered by Ramsey (1928)), but there is no practical sense in supposing that *inherited* state variables can actually change their distribution in response to *subsequent* choice. The consequence of the unconditional optimisation approach is a first-order condition equivalent to (2.40), only with $\beta = 1$.

Our method therefore shares the discounting structure of conventional methods (including the timeless perspective approach) when applied to endogenous state variables; but deviates from it where variables feature in expectational constraints. The former follows, as usual, from the fact that today's past cannot be changed; but nor, any longer, can yesterday's vision of the future – and it is this that accounts for the latter.

[72] We will encounter this problem in more detail in the second and third of the examples below.

3 Example: A deterministic inflation bias model

To demonstrate the implications of the veil of ignorance policymaking approach we first draw upon a simple linear-quadratic example that has been widely studied in the New Keynesian literature, but nonetheless gives insight into our method: a purely deterministic problem in which the policymaker faces a choice over inflation rates and the output gap, subject to a linear 'New Keynesian Phillips Curve'. Market power on the part of firms constrains the natural level of output below its efficient level, and gives an incentive to inflate the economy, even at the cost of some inflation (which induces inefficient price dispersion in the event that pricing is modelled in the spirit of Calvo (1983)). Details of the underlying non-linear model are widely known – see, for instance, Woodford (2003).[73] Despite the familiarity of the model, we retain the more involved notational conventions from above, so as to focus attention on the methodological steps taken.

3.1 The unrestricted problem: one-off versus repeated choice

We first outline the problem when all future state- and date-contingent outcomes can be chosen in an initial time period, time 0. Absent any informational constraints the policymaker chooses (history- and time-contingent) values for inflation, $\pi(h,t)$, and the output gap, $y(h,t)$, from time 0 onwards to maximise W_0:

$$W_0 = -\frac{1}{2}\sum_{t=0}^{\infty}\beta^t \int_{H^\infty} \left[\pi(h,t)^2 + \omega\left(y(h,t) - y^*\right)^2\right] dF_t\left(h|h^0\right) \tag{3.1}$$

This choice is subject to a 'Phillips curve' relationship that must hold for all (h,t) pairs:

$$\pi(h,t) = \beta \int_{H^\infty} \pi\left(h',t+1\right) dF_1\left(h'|h\right) + \gamma y(h,t) \tag{3.2}$$

[73] As is well known, the setup is obtained by taking a linear approximation to the basic New Keynesian model around a zero-inflation steady state. This renders the exercise that follows slightly forced, since we seek an optimal constant inflation rate in the linear-quadratic problem – and this rate may not in fact be zero. The example's simplicity nonetheless makes it useful in expositing the general optimality procedure.

In general the triple $(h_\varepsilon, z_{-1}, t)$ is sufficiently fine information to characterise optimal choice from the perspective of time 0 in response to any (h, t) pair – since other elements of history are either irrelevant to the objective and constraint set at time 0 (past variables not in z_{-1}) or can themselves be determined from the same information. Since this model features no state variables and no exogenous shocks, this in turn implies optimal time-0 choice can be written as a function of t alone: we drop h from the notation hereafter.

The first-order conditions characterising optimal choice of $\pi(t)$ and $y(t)$, using multipliers $\lambda(t)$ on the constraint (3.2), are:

$$\pi(0) - \lambda(0) = 0 \tag{3.3}$$

$$\pi(t) - \lambda(t) + \lambda(t-1) = 0 \tag{3.4}$$

for $t > 0$ and

$$\omega(y(t) - y^*) + \gamma\lambda(t) = 0 \tag{3.5}$$

for $t \geq 0$.

With some fairly trivial algebra these conditions (together with (3.2)) solve to give the *time*-0 *optimum*:

$$\pi(t) = \frac{\omega}{\gamma}(1-\phi)\, y^* \phi^t \tag{3.6}$$

$$y(t) = y^* \phi^{t+1} \tag{3.7}$$

where we have defined ϕ as the (stable) root:

$$\phi \equiv \frac{1 + \beta + \frac{\gamma^2}{\omega} - \sqrt{\left(1 + \beta + \frac{\gamma^2}{\omega}\right)^2 - 4\beta}}{2\beta}$$

The optimal solution from the perspective of time 0 thus involves a substantial initial inflation, pulling output closer to its efficient level y^*, followed by a gradual decay back to zero values for both choice variables. The time-inconsistency in this solution is clear: policymakers in periods subsequent to time 0 would like to implement precisely the same initial inflation – knowing that past expectations can no longer be affected detrimentally by their so doing. The solution can only obtain if choice over all variables at all time periods is the responsibility of the single policymaker who exists at the start of time.

Optimal policy from a timeless perspective follows directly from the time-0 optimum. Recall that this approach again assumes policy can be chosen for all future periods at the start of time, but does so under the assumption that choices should be best from the perspective of a time period in the distant past. In this context, the implication is that we should implement in every period the values $\pi(\infty)$ and $y(\infty)$ defined under (3.6) and (3.7). Hence we have the *timeless perspective solution*:

$$\pi(t) = 0 \tag{3.8}$$

$$y(t) = 0 \tag{3.9}$$

If instead we assume that choice over variables in any given time period is the responsibility of the policymaker who is alive in that period, maximising W_t with respect to $(\pi(t), y(t))$ alone, we can find a Stackelberg ('discretionary') equilibrium, in which choice is always made subject to the best response functions of future decision-makers (together with private-sector expectations, which must be consistent with equilibrium play after all histories). Restricting attention to Markov-stationary equilibria – which in this context implies constant values for the choice variables, since we have a stationary problem with no shocks – and denoting by $\lambda(t)$ the multiplier on (3.2) as before, we have first-order conditions for every $t \geq 0$:

$$\pi(t) - \lambda(t) = 0 \tag{3.10}$$

$$\omega(y(t) - y^*) + \gamma\lambda(t) = 0 \tag{3.11}$$

These hold together with (3.2), yielding a unique Markov-stationary equilibrium solution for all $t \geq 0$, the *discretionary solution*:

$$\pi(t) = \frac{\gamma}{1-\beta+\frac{\gamma^2}{\omega}} y^* \tag{3.12}$$

$$y(t) = \frac{1-\beta}{1-\beta+\frac{\gamma^2}{\omega}} y^* \tag{3.13}$$

In general these will not be Pareto efficient outcomes for the game among the policymakers: different choices by each could improve outcomes for all. The framework in this paper is geared towards obtaining these Paretian gains, without doing so by tailoring outcomes to the particular preferences of one policymaker alone. This specific example clearly derives added practical relevance from the common practice of devolving interest-rate-setting to independent central banks, as a way to institutionalise the gains from moving away from discretion.

3.2 Optimal policy from behind a veil of ignorance

We now assume choice must be made under the veil of ignorance framework outlined above. Since there are no lagged endogenous variables featuring in the model's structural equations and no exogenous shocks, the sets E^∞ and Z are redundant – so we need only choose a single pair of values (π, y) to maximise the objective:

$$W_0 = -\frac{1}{2}\sum_{t=0}^{\infty}\beta^t\left[\pi^2 + \omega\left(y-y^*\right)^2\right] \tag{3.14}$$

subject to:

$$\pi = \beta\pi + \gamma y \tag{3.15}$$

Denoting by λ the multiplier on (3.15), we have the following first-order conditions:

$$\pi - (1-\beta)\lambda = 0 \tag{3.16}$$

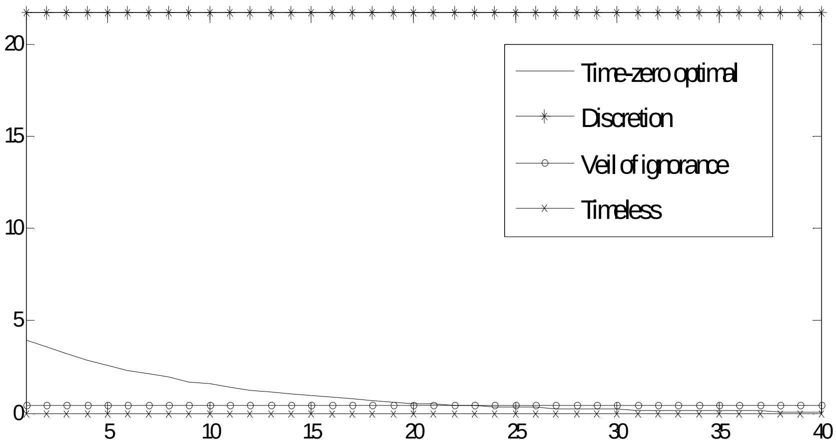

Figure 3.1: Response of inflation under different policy approaches

$$\omega\left(y-y^{*}\right)+\gamma\lambda=0 \tag{3.17}$$

We then obtain the *veil of ignorance solution*:

$$\pi=\frac{(1-\beta)\,\gamma}{(1-\beta)^{2}+\frac{\gamma^{2}}{\omega}}y^{*} \tag{3.18}$$

$$y=\frac{(1-\beta)^{2}}{(1-\beta)^{2}+\frac{\gamma^{2}}{\omega}}y^{*} \tag{3.19}$$

This implies a positive level of inflation that is nonetheless strictly lower than the 'discretionary' outcome. Note that this contrasts with the 'timeless perspective' outcome, which chooses optimally from the perspective of a time period in the distant past – which in this case (by equation (3.6)) implies setting inflation and output to zero. Figure 3.1 shows the paths for inflation and output over time under the four policy approaches: time-zero optimisation, Markov-stationary 'discretion', timeless perspective optimisation, and veil of ignorance optimisation.[74]

[74] For this we take parameter values from Woodford (2003), where a similar policy comparison is made: $y^{*}=0.2$, $\beta=0.99$, $\gamma=0.024$ and $\omega=0.048$.

We can see here the consequences of adopting the veil of ignorance perspective: inflation is reduced well below the value implied by discretionary policy, but remains slightly positive – at a value (roughly 0.4 per cent) that is surprisingly close to the *quarterly* rate targeted by many central banks. The time-0 optimal policy, on the other hand, demands a large initial inflation followed by gradual disinflation – tending towards zero. The timeless perspective demands zero inflation in all time periods. The key difference between time-0 optimality and optimality from behind a veil of ignorance can be seen by contrasting the two first-order conditions (3.4) and (3.16). As reflected in the dating of the multipliers, the former is taken under the assumption that any change to the (planned) inflation rate in time t influences the policy options available at time $t-1$ (for $t>0$). Shadow costs due to this effect are weighted at rate β^{-1} relative to direct effects on time-t loss – resulting in the form given.[75] At time 0 there are no expectational effects of choice. But if we constrain the policymaker to choose a single rate of inflation for all time periods then changes to this rate at the margin affect expectations in *all* time periods from 0 onwards, as well as having direct effects on within-period welfare in the *same* set of time periods. Thus the relative weighting of expectational effects by β^{-1} is no longer appropriate, which accounts for the enduring presence of β in equation (3.16) – and thus, by comparison with the timeless perspective result, the positive inflation rate.

The timeless perspective method adopts a different normative approach from ours, asking not what is best from the perspective of a policymaker denied time-specific information, but rather what is best from the perspective of a hypothetical policymaker choosing under full information in the distant past. This example highlights an uncomfortable implication of the timeless perspective approach: in the relatively simple class of deterministic choice problems without 'backward-looking' constraints it does not select the best possible time-invariant values for the choice variables. In this specific case, when choosing under the timeless perspective one will tend to exaggerate the benefits of keeping inflation low in order to relax earlier expectational constraints – *even when those constraints never in fact*

[75] The β in the Phillips curve cancels with this β^{-1} when deriving (3.4).

constrained choice from time 0 onwards. Under the time-0 optimal path the policymaker exploits the relatively benign expectations present at time 0 to inflate substantially and raise output close to y^* – knowing that there are no expectational effects of doing so. Subsequent disinflation is, in a sense, the price of this. The timeless perspective approach pays the price without enjoying the benefit – the latter hypothetically reserved for long-dead ancestors.

Confirming Pareto dominance is a matter of simple algebra. Under the timeless perspective approach the value of the objective function is given be W_0^{TP}:

$$W_0^{TP} = -\frac{1}{2}\frac{1}{1-\beta}\omega\left(y^*\right)^2 \tag{3.20}$$

The veil of ignorance method instead yields a value W_0^{VI}:

$$\begin{aligned} W_0^{VI} &= -\frac{1}{2}\frac{1}{1-\beta}\left[\left(\frac{(1-\beta)\gamma}{(1-\beta)^2+\frac{\gamma^2}{\omega}}y^*\right)^2 + \omega\left(\frac{\frac{\gamma^2}{\omega}}{(1-\beta)^2+\frac{\gamma^2}{\omega}}y^*\right)^2\right] \\ &= -\frac{1}{2}\frac{1}{1-\beta}\frac{\gamma^2}{\omega(1-\beta)^2+\gamma^2}\omega\left(y^*\right)^2 \\ &> W_0^{TP} \end{aligned} \tag{3.21}$$

Of course, there is no ‘correct’ normative approach to solving problems with time inconsistency; we just have a range of methods across which the information structure and set of priorities that generate choice in each time period are varied. But the fact that timeless perspective policies are Pareto dominated within this relatively simple (deterministic, purely forward-looking) class of problems does little to enhance that method’s normative appeal. Moreover, there is no obvious compensating advantage to the timeless perspective approach. For instance even when one moves away from purely *deterministic* forward-looking models to a simple stochastic ‘stabilisation bias’ problem, Blake (2001) has shown that the timeless perspective approach is dominated in expectation by a rule that coincides with our veil of ignorance solution.

4 Example: Optimal redistribution with participation constraints

The previous example was aimed chiefly at highlighting the difference between optimality from behind a veil of ignorance and alternative solution approaches, in the simplest model possible. The next makes far fuller use of the general theoretical framework that we have set out. It is a simple model of redistributive social insurance, in which a utilitarian policymaker must allocate consumption knowing that some agents may opt to leave the scheme if they consider that they would be better off consuming under 'autarky'. We will illustrate a number of appealing general theoretical properties of the veil of ignorance solution in the first sub-section, before turning to a specific computed example.

4.1 The general model

The basic problem is adapted from the works of Marcet and Marimon (1992) and Kocherlakota (1996). There are N infinitely-lived agents in an economy, each of whom is endowed with a stochastic income stream through time,[76] with the income of agent n at any given point in time denoted $y^n \in \Upsilon$.[77] A social planner has weighted-utilitarian preferences across the subjective welfare of these agents, with an objective criterion at time 0 given by:

$$W_0 = \sum_{n=1}^{N} \alpha^n U_0^n \tag{4.1}$$

where U_0^n is agent n's subjective welfare assessed at time 0 – a standard function of n's consumption alone:

$$U_0^n = \sum_{t=0}^{\infty} \beta^t \int_{H^\infty} u^n \left(c^n \left(h, t\right)\right) dF_t \left(h | h^0\right) \tag{4.2}$$

[76] It will be useful in specific examples to assume an infinite number of agents existing on a continuum as opposed to a finite N.

[77] We choose not to index y^n by time, to retain consistency with the general presentation. The same applies to the stochastic real interest rate introduced below.

As above, $F_t\left(h|h^0\right)$ is an equilibrium-consistent distribution across histories at time t, conditional upon observing h^0 at time 0. $c^n(h,t): H^\infty \times \mathbb{Z} \to \mathbb{R}_+$ is agent n's consumption at time t following history h. $u^n(\cdot): \mathbb{R}_+ \to \mathbb{R}$ is a within-period utility function that is continuously differentiable and concave.

The planner must allocate income to the different agents subject to an aggregate resource constraint holding at all t:

$$\sum_{n=1}^{N} y^n + R'B\left(h\backslash 1, t-1\right) \geq \sum_{n=1}^{N} c^n(h,t) + B(h,t) \tag{4.3}$$

where $R' \in \mathbb{R}_{++}$ is an exogenous, stochastic real interest rate realised at time $t-1$ (that is, the penultimate entry for the relevant series in h), and $B(h,t): H^\infty \times \mathbb{Z} \to \mathbb{R}$ is the planner's quantity of savings following history h at time t. In what follows we denote the *contemporary* real interest rate – to be paid out on current bond acquisitions – simply by R (this is the terminal entry for the relevant series in h). The prior condition must hold together with a relevant transversality ('no Ponzi') restriction, ensuring dynamic solvency.

In addition, each agent has the option of migrating away from the economy permanently, to enjoy their endowment income in solitude. To prevent this, the planner must ensure in each period that every agent anticipates a higher utility from staying put than from migrating. The participation constraint for each n and all h, t is thus:

$$U_t^n \geq \sum_{s=t}^{\infty} \beta^{s-t} \int_{\Upsilon} u^n(y^n)\, dF_{s-t}(y^n|h) \tag{4.4}$$

For the problem not to be trivial we must assume that these participation constraints will bind for some agents in one or more periods (under histories that are observed with strictly positive probability). But the constraints will imply time-inconsistency. The best level of utility to promise, at time t, to give to agent n at time s is likely to differ from the best level of utility to give agent n from the perspective of time s – since the t-dated participation constraint will by then have ceased to be a concern.

4.2 The time-zero optimum

We first sketch the basic characteristics of the optimal solution from the perspective of time 0. Attaching present-value multiplier $\lambda_R(h,t)$ to the resource constraint (4.3), and $\lambda_P^n(h,t)$ to the participation constraint (4.4), we have as a first-order condition on optimal choice of $c^n(h,t)$:

$$\left[\alpha^n + \sum_{s=0}^{t} \lambda_P^n(h\backslash(t-s), s)\right] \frac{\partial u^n(c^n(h,t))}{\partial c^n(h,t)} - \lambda_R(h,t) = 0 \tag{4.5}$$

This implies that the ratio of marginal utilities between agents m and n evolves over time according to the relationship:

$$\frac{\frac{\partial u^n(c^n(h,t))}{\partial c^n(h,t)}}{\frac{\partial u^m(c^m(h,t))}{\partial c^m(h,t)}} = \frac{\alpha^m + \mu^m(h,t)}{\alpha^n + \mu^n(h,t)} \tag{4.6}$$

where we have defined the non-stationary, agent-specific variable $\mu^n(h,t)$ by:

$$\begin{aligned} \mu^n(h,t) &\equiv \sum_{s=0}^{t} \lambda_P^n(h\backslash(t-s), s) \\ &= \mu^n(h\backslash 1, t-1) + \lambda_P^n(h,t) \end{aligned} \tag{4.7}$$

(the last line here for $t > 0$ only).

This result highlights again the uncomfortable consequences that can follow from imposing the time-0 optimum. Note that $\lambda_P^n(h,t) \geq 0$ will always hold, so μ^n is non-decreasing through time. If there were a non-zero probability at every horizon of participation constraints binding then the ratio of α^n to μ^n could even become negligible as time progresses. That is, the 'fundamental' concern the policymaker has for the utility of any given agent becomes dominated by a need to abide by past commitments.[78] If we assume a symmetric objective, so $\alpha^n = \alpha^m$ for all $n, m \leq N$, this implies future

[78] There is a limit to this effect in versions of the model where one makes the twin simplifying assumptions that the real interest rate is constant and equal to β^{-1}, and Υ is a discrete set with Markov transition probabilities for individual income draws. In that situation all agents who have received the highest possible income draw at *some* point in time since period 0 will be afforded a weight $\overline{\mu}$ that is permanently high but will increase no further. To this $\overline{\mu}$ is associated

policymakers potentially being asked to accept extreme inequality as the price of past incentives – despite having a basic preference structure altogether more egalitarian. It is hard to see how this 'tyranny of the past' could be sustained in a democratic society.

The non-stationarity of the μ^n weights derives from the time-specific assessment of marginal gains and losses. By definition, any alterations to the function $c^n(h,t)$ are known to be particular to time t. Any effects on incentives s periods prior to these changes – due to the associated effects on the within-period utility function at time t – are therefore relevant to policy options at time $t-s$. This implies weighting those 'incentive effects' at rate β^{-s} relative to the direct effects on utility at time t of changing $c^n(h,t)$, given the time preference structure contained in W_0. But the incentive effects in turn are weighted by a coefficient β^s when considering the impact of a change in period-t utility on the lifetime utility for an agent alive at time $t-s$. Thus the effects of time preference cancel: changes to $c^n(h,t)$ can be used through an incentive channel to improve outcomes at $t-s$, which matter more to the policymaker than those at t, but *strength* of this incentive channel is decaying at a rate that just offsets the relative time preference. Again, this implies that if one weren't able to associate marginal changes to consumption with specific time periods the structure of the optimal solution could change.

The first-order condition on $B(h,t)$ requires:

$$-\lambda_R(h,t)+\beta R\int_{H^\infty}\lambda_R\left(h',t+1\right)dF_1\left(h'|h\right)=0 \tag{4.8}$$

This permits an intertemporal optimality condition to be written for the consumption path of each agent n:

$$\begin{aligned}&(\alpha^n+\mu^n(h,t))\frac{\partial u^n(c^n(h,t))}{\partial c^n(h,t)}\\&=\beta R\int_{H^\infty}\left(\alpha^n+\mu^n\left(h',t+1\right)\right)\frac{\partial u^n(c^n(h',t+1))}{\partial c^n(h',t+1)}dF_1\left(h'|h\right)\end{aligned} \tag{4.9}$$

a *permanently* high, constant income level, sufficient to incentivise participation even were the agent again to draw the maximal element in Υ.

Since $(\alpha^n + \mu^n(h', t+1)) \geq (\alpha^n + \mu^n(h,t))$ for all h', this implies a distortion to the standard intertemporal Euler condition, so long as incentive constraints bind with non-negligible probability for agent n at time $t+1$. The marginal benefits to the policymaker of transferring consumption to agent n at time $t+1$ include gains from relaxing agent n's participation constraint in that time period – a constraint that is not present at t. If a fully binding commitment not to migrate could be made *ex ante* by the agent, incentive compatibility constraints would never apply beyond the initial time period. The solution would then see no change in $\mu^n(h', t)$ through time – and we would have greater frontloading of consumption for any given interest rate path.

4.3 Optimality from behind a veil of ignorance

We now consider the implications of imposing the 'veil of ignorance' approach to optimality within this framework. The model features one lagged endogenous variable, B, together with an $(N+1)$-dimensional vector of exogenous stochastic variables in each time period, comprised of $\{y^n\}_{n=1}^N$ and R (collectively denoted $\varepsilon \in E$ as before, where the set E is here given by $\Upsilon^N \times \mathbb{R}_{++}$, and with associated exogenous history $h_\varepsilon \in E^\infty$). We denote the inherited stock of savings $B_{-1} \in \mathbb{R}$. For reasons discussed above, since we have both forward- and backward-looking constraints we must divide up choice into forward- and backward-looking problems, before mapping between the two to obtain a solution. Consistent with our general discussion, the forward-looking policymaker chooses functions $\{c^n(h_\varepsilon, B)\}_{n=1}^N$, with $c^n(h_\varepsilon, B) : E^\infty \times \mathbb{R} \to \mathbb{R}_+$ for all n, to solve:

$$\max_{\{c^n(h_\varepsilon,B)\}_{n=1}^N} \frac{1}{1-\beta} \int_{E^\infty \times \mathbb{R}} \left\{ \sum_{n=1}^N \alpha^n u^n\left(c^n(h_\varepsilon, B)\right) \right. \tag{4.10}$$
$$+ \sum_{n=1}^N \sum_{s=0}^\infty \lambda_P^{n,s}(h_\varepsilon, B)\, \beta^s \left[u^n\left(c^n(h_\varepsilon, B)\right) - u^n(y^n)\right]$$
$$\left. + \lambda_R(h_\varepsilon, B) \left[\sum_{n=1}^N y^n + R'B - \sum_{n=1}^N c^n(h_\varepsilon, B) - B(h_\varepsilon, B) \right] \right\} dF^0(h_\varepsilon, B)$$

This is done treating the functions $B(h_\varepsilon, B)$, $\lambda_P^{n,s}(h_\varepsilon, B)$ and $\lambda_R(h_\varepsilon, B)$ as fixed, and for a distribution F^0 with full support on $E^\infty \times \mathbb{R}$.[79] Note that the coefficient β^s in the second line here comes directly from the participation constraint.[80] The backward-looking policymaker chooses values of $B(h_\varepsilon, B_{-1}, t) : E^\infty \times \mathbb{R} \times \mathbb{Z} \to \mathbb{R}$ to solve:

$$
\begin{aligned}
&\max_{B(h_\varepsilon, B_{-1}, t)} \sum_{t=0}^{\infty} \beta^t \int_{E^\infty} \left\{ \sum_{n=1}^{N} \alpha^n u^n \left(c^n \left(h_\varepsilon, B_{-1}, t\right)\right) \right. \\
&+ \sum_{n=1}^{N} \lambda_P^n \left(h_\varepsilon, B_{-1}, t\right) \sum_{s=0}^{\infty} \beta^s \int_{E^\infty} \left[u^n \left(c^n \left(h'_\varepsilon, B_{-1}, t+s\right)\right) - u^n \left(y^{n\prime}\right)\right] dF_s \left(h'_\varepsilon | h_\varepsilon\right) \\
&+ \lambda_R \left(h_\varepsilon, B_{-1}, t\right) \left[\sum_{n=1}^{N} y^n + R' B \left(h_\varepsilon \backslash 1, B_{-1}, t-1\right) \right. \\
&\left. \left. - \sum_{n=1}^{N} c^n \left(h_\varepsilon, B_{-1}, t\right) - B \left(h_\varepsilon, B_{-1}, t\right) \right] \right\} dF_t \left(h_\varepsilon | h_\varepsilon^0\right)
\end{aligned} \tag{4.11}
$$

Similarly, this problem is solved for given $c^n(h_\varepsilon, B_{-1}, t)$, $\lambda_P^n(h_\varepsilon, B_{-1}, t)$ and $\lambda_R(h_\varepsilon, B_{-1}, t)$ functions. The multiplier functions are then chosen in a manner that ensures complementary slackness holds.

Any 'veil of ignorance' solution as defined above must satisfy a first-order condition with respect to choice of $c^n(h_\varepsilon, B)$:

$$
\left[\alpha^n + \sum_{s=0}^{\infty} \beta^s \lambda_P^{n,s}(h_\varepsilon, B)\right] \frac{\partial u^n \left(c^n \left(h_\varepsilon, B\right)\right)}{\partial c^n \left(h_\varepsilon, B\right)} - \lambda_R \left(h_\varepsilon, B\right) = 0 \tag{4.12}
$$

Optimal choice of $B(h_\varepsilon, B_{-1}, t)$, meanwhile, requires generically:

$$
-\lambda_R \left(h_\varepsilon, B_{-1}, t\right) + \beta R \int_{E^\infty} \lambda_R \left(h'_\varepsilon, B_{-1}, t+1\right) dF_1 \left(h'_\varepsilon | h_\varepsilon\right) = 0 \tag{4.13}
$$

At any veil of ignorance optimum, this latter condition implies the multipliers $\lambda_R(h_\varepsilon, B)$ defined on

[79] R' is again used here to denote the penultimate entry for the real interest rate in h_ε. The last entry (which is of relevance to dynamic choice) will again be expressed simply as R.

[80] In terms of our general notation, the function g_s here is given by $\beta^s \left[u^n \left(c^n \left(h_\varepsilon, B\right)\right) - u^n \left(y^n\right)\right]$.

the reduced space $E^{\infty} \times \mathbb{R}$ must satisfy:

$$-\lambda_R(h_\varepsilon, B) + \beta R \int_{E^\infty} \lambda_R\left(h'_\varepsilon, B(h_\varepsilon, B)\right) dF_1\left(h'_\varepsilon | h_\varepsilon\right) = 0 \tag{4.14}$$

Comparing equations (4.5) and (4.8) with (4.12) and (4.14), the effects of the 'veil of ignorance' can be seen. In particular, the choice of savings depends upon essentially the same factors as before: the relative shadow value of relaxing the resource constraint in one time period against the next. Since changes to 'backward-looking' policy can never influence B_{-1}, it is appropriate even under the veil of ignorance setup to discount the later effects here at rate β.

But the choice of $c^n(h_\varepsilon, B)$ differs importantly across the two approaches. The forward-looking policymaker has been constrained to believe that the chance of any pair $(h_\varepsilon, B) \in E^{\infty} \times \mathbb{R}$ characterising any particular time period is independent of the precise time period in question. Thus for any t there is an equal chance that welfare in that period will be influenced by choice of $c^n(h_\varepsilon, B)$ directly, or through an incentive channel operating at an arbitrary horizon s. The latter must be discounted at rate β^s, in accordance with the impact that within-period outcomes at time $t+s$ have on lifetime utility at time t. But there is no offsetting β^{-s} weighting in this case: both the direct and the incentive effects are experienced at time t with identical probability. Moreover, there is no limit to the horizon at which changes to $c^n(h_\varepsilon, B)$ influence utility in any given time period. This is why the sum in (4.12) has no upper limit. By contrast, when policy is time-specific the choice in period t can only have an impact on the incentive compatibility constraints for periods 0 through to t.

The consequences of this differential treatment of forward-looking effects can be seen in equations (4.15) and (4.16) – the 'veil of ignorance' equivalents of (4.6) and (4.7):

$$\frac{\frac{\partial u^n(c^n(h_\varepsilon,B))}{\partial c^n(h_\varepsilon,B)}}{\frac{\partial u^m(c^m(h_\varepsilon,B))}{\partial c^m(h_\varepsilon,B)}} = \frac{\alpha^m + \mu^m(h_\varepsilon, B)}{\alpha^n + \mu^n(h_\varepsilon, B)} \tag{4.15}$$

$$\mu^n (h_\varepsilon, B) \equiv \sum_{s=0}^{\infty} \beta^s \lambda_P^{n,s} (h_\varepsilon, B) \tag{4.16}$$
$$= \beta \sum_{B' \in \mathbb{R}} \mu^n \left(h_\varepsilon \backslash 1, B'\right) P_{-1}^z \left(B' | h_\varepsilon \backslash 1, h_\varepsilon, B\right) + \lambda_P^{n,0} (h_\varepsilon, B)$$

Assuming, realistically, that the mapping from past to current choices of B will be one-to-one (for any given h_ε) we can denote by $B^{-1} (h_\varepsilon, B) : E^\infty \times \mathbb{R} \to \mathbb{R}$ the unique value of $B' \in \mathbb{R}$ such that $P_{-1}^z (B' | h_\varepsilon \backslash 1, h_\varepsilon, B) = 1$. Then the last expression can be written in simpler form:

$$\mu^n (h_\varepsilon, B) \equiv \beta \mu^n \left(h_\varepsilon \backslash 1, B^{-1} (h_\varepsilon, B)\right) + \lambda_P^{n,0} (h_\varepsilon, B)$$

As can be seen, the μ^n terms are no longer non-decreasing as we progress through history: instead they converge to zero at rate β under all histories for which incentive compatibility constraints do not bind. This has important consequences for (4.15): the underlying α^n weights will continue to have a non-negligible impact on the ratio of marginal utilities between any two agents as time progresses, even when incentive compatibility constraints bind with non-negligible probability for all agents at all horizons. So instead of priorities in the longer term being driven entirely by a desire to satisfy past constraints, the 'true' social preference structure retains permanent importance. This seems a substantial advantage if one wishes for a choice mechanism that will permit future utility to be used in current incentive packages, whilst at the same time standing a strong chance of enduring future democratic choice.

The equivalent to condition (4.9) in the 'veil of ignorance' case is:

$$(\alpha^n + \mu^n (h_\varepsilon, B)) \frac{\partial u^n (c^n (h_\varepsilon, B))}{\partial c^n (h_\varepsilon, B)} \tag{4.17}$$
$$= \beta R \int_{E^\infty} \left(\alpha^n + \mu^n \left(h'_\varepsilon, B (h_\varepsilon, B)\right)\right) \frac{\partial u^n (c^n (h'_\varepsilon, B (h_\varepsilon, B)))}{\partial c^n (h'_\varepsilon, B (h_\varepsilon, B))} dF_1 \left(h'_\varepsilon | h_\varepsilon\right)$$

Since the $\mu^n (h_\varepsilon, B)$ terms are no longer non-decreasing, we can no longer say how the conventional

Euler condition is being distorted. In particular, if the contemporary value of y^n contained in h_ε is notably high, one might well expect $\mu^n(h_\varepsilon, B)$ to exceed its successors on average, so that the marginal utility of consumption for agent n after history (h_ε, B) is lower than its expected value across successor histories (weighted by βR). So the provision of incentives is now being 'frontloaded' relative to the full information case: the policymaker behind a veil of ignorance is content to let bygones fade. Again, this follows from the increasingly weak impact that outcomes at $t+s$ could have on incentives at t as s becomes large – coupled with the fact that behind a veil of ignorance one must rank these incentive effects 'equally' against the possibility of direct utility gains at time t.

This example suggests that the qualitative consequences of adopting the veil of ignorance optimality perspective could be substantial. Woodford (2010) rightly argues that the difference between the 'unconditionally optimal' and 'timeless perspective' approaches is likely to be very small in a model for which T, the maximum horizon at which expectations matter, is equal to 1. He makes the point with reference to a stabilisation bias problem that is essentially a stochastic version of our first example (with y^* additionally set to zero) – and in which the unconditionally optimal and veil of ignorance methods will coincide, due to the absence of endogenous state variables. The point is that the difference only hinges on the inclusion or non-inclusion of a single β^{-1} term as a coefficient in the relevant first-order condition. If β is close to 1 in any event, this will have little effect.

But if T is substantially larger than 1 – and it is infinite in the example here – the argument no longer applies. Clearly β^{-T} is unbounded as T grows, and the difference between weighting by it for pure time preference (as when the timing of marginal policy changes is known) and neglecting to do so (as under the veil of ignorance approach) makes the difference here between stationarity and non-stationarity in the policymaker's utility weighting.

The example also draws attention to an important technical advantage that the veil of ignorance approach has over the timeless perspective: in the event that T is infinite the timeless perspective policy will not in general be well defined. Time-0 optimal policy is non-stationary, and we do not necessarily have convergence on any particular set of μ^n weights as the time elapsed since optimisation grows. Even if we were to see this convergence, it is unlikely that one could *impose* such values from the start, whilst still allowing all participation and resource constraints to bind – a point that becomes clearer in the specific numerical example that follows. So it may not even be *feasible* in this model to advocate a policy that would have been best from a point in the distant past, not least because a policymaker in the distant past might well have chosen to build up a substantial asset stock as the compound weights on agents' within-period utilities accumulated. The veil of ignorance method, by contrast, induces stationary representations for the μ^n weights, and thereby easily extends to this setting in a way that the timeless perspective does not.

4.4 A specific case

To give more content to the analytical results just derived we now consider a specific case of the preceeding model, using a number of simplifying assumptions that allow for an intuitive solution approach. First, we assume that a continuum of households exists – indexed by a position on the unit interval. The policymaker has symmetric preferences across individuals, weighting each household equally in a social welfare function that therefore takes the following form:

$$W_0 = \int_0^1 U_0^i di \tag{4.18}$$

Household utility functions take the familiar CES form:

$$U_0^i = \sum_{t=0}^{\infty} \beta^t \int \frac{c^i(h,t)^{1-\sigma} - 1}{1-\sigma} dF_t\left(h|h^0\right) \tag{4.19}$$

and endowment income processes are assumed for now to be iid across households and time, with household i's income in any given time period, y^i, given by:

$$y^i = \begin{cases} y^h \text{ with probability } p^h \in (0,1) \\ y^l \text{ with probability } 1 - p^h \end{cases} \tag{4.20}$$

with $y^h > y^l$.

Any solution to the redistribution problem must therefore satisfy the incentive compatibility constraint:

$$U_t^i \geq \begin{cases} \frac{(y^h)^{1-\sigma}-1}{1-\sigma} + \frac{\beta}{1-\beta}\left[p^h \frac{(y^h)^{1-\sigma}-1}{1-\sigma} + \left(1-p^h\right)\frac{(y^l)^{1-\sigma}-1}{1-\sigma}\right] & \text{if } y_t^i = y^h \\ \frac{(y^l)^{1-\sigma}-1}{1-\sigma} + \frac{\beta}{1-\beta}\left[p^h \frac{(y^h)^{1-\sigma}-1}{1-\sigma} + \left(1-p^h\right)\frac{(y^l)^{1-\sigma}-1}{1-\sigma}\right] & \text{if } y_t^i = y^l \end{cases} \tag{4.21}$$

at every horizon t.

Finally, we assume that the gross real interest rate R is constant through time, and equal to the inverse of the discount factor β. Veil of ignorance first-order conditions with respect to $c^i(h_\varepsilon, B)$ and $B(h_\varepsilon, B)$ respectively (the latter having been mapped onto the space $E^\infty \times \mathbb{R}$) now take the form:

$$\left[1 + \sum_{s=0}^{\infty} \beta^s \lambda_P^{i,s}(h_\varepsilon, B)\right] c^i(h_\varepsilon, B)^{-\sigma} - \lambda_R(h_\varepsilon, B) = 0 \tag{4.22}$$

and

$$-\lambda_R(h_\varepsilon, B) + \int \lambda_R\left(h'_\varepsilon, B(h_\varepsilon, B)\right) dF_1\left(h'_\varepsilon | h_\varepsilon\right) = 0 \tag{4.23}$$

for all (h_ε, B) pairs, with the multipliers on participation and resource constraints defined as before.

4.4.1 Solution approach

Our assumption of a continuum of households receiving symmetric welfare weightings, together with the constancy of R, ensures, by a standard law of large numbers, that the history of exogenous variables h_ε will be of no relevance to the policymaker's optimal choice of B: any two realisations of h_ε will almost surely be equivalent from the policymaker's perspective (though individual-specific income histories will of course still vary). Hence the shadow value of relaxing the aggregate resource constraint will be a function of inherited assets alone, and we may denote it $\lambda_R(B)$. Equation (4.23) then implies that $\lambda_R(B) = \lambda_R(B(B))$, and since the shadow value of additional resources must be decreasing in the existing asset stock we must have that $B(B) = B$: the interest rate of β^{-1} is just sufficient to keep wealth constant. Absent dependence of λ_R on h_ε we additionally have no reason for $c^i(h_\varepsilon, B)$ to depend on income levels of agents other than i, and we denote the history of i-specific income draws h^i_ε.

To obtain a solution we can further note that partcipation constraints will never have been binding for the most unlucky individual – that is, one who has received income y^l in every period in history. Otherwise, given the policymaker's natural preference for greater equality under diminishing marginal utility, transfers away from an agent for whom participation constraints are satisfied with *strict* equality could deliver an improvement. Denoting this agent's history $h^i_{\underline{\varepsilon}}$, we have from (4.22) that the shadow value of income in each time period must equal the most unlucky agent's marginal utility of consumption – which in turn must be constant. Since agents receiving an income level of y^l some time after previously receiving a high income y^h must receive at least as much utility as those who have experienced y^l forever, we can then solve for the marginal utility of consumption for other agents by conjecturing that (contemporary) participation constraints will only bind for agents with a current income of y^h, and that the solution will ensure all individuals receiving y^h will experience the *same* levels of consumption and expected utility, irrespective of their specific prior histories (given the decaying nature of the multipliers as high income shocks recede, it is as if their Pareto weights must be 'topped up' to a maximum level to ensure participation each time autarky utility is at its highest).[81] Thus there will exist some $\overline{\mu}$ such that whenever h_ε specifies an income of y^h for agent i in the most recent time period we must have:

$$\overline{\mu} = \sum_{s=0}^{\infty} \beta^s \lambda_P^{i,s}(h_\varepsilon, B) \tag{4.24}$$

[81] In a similar model Kocherlakota (1996) refers to this as an 'amnesia' property: when constraints bind the past is completely forgotten.

In general one would expect $\overline{\mu}$ to be higher the lower is B, since a policymaker with access to a large quantity of assets will be able to engineer a more equal distribution of resources without concern for participation constraints (for a high enough stock of assets participation constraints will not bind for any agent, and the outcome will be complete equality). For an agent i who last received a high income r periods ago, condition (4.22) then becomes:

$$[1+\beta^r\overline{\mu}]\,c^i\left(h_\varepsilon,B\right)^{-\sigma}-\lambda_R\left(h_\varepsilon,B\right)=0 \tag{4.25}$$

We simplify notation now, and denote the consumption of an agent who last received a high income shock r periods ago $c(r)$; since assets are being held constant through time, this value will be fixed for a given initial asset stock B. We can express $c(r)$ as a function of $c(\infty)$:

$$c\left(r\right)=\left(1+\beta^r\overline{\mu}\right)^{\frac{1}{\sigma}}c\left(\infty\right) \tag{4.26}$$

Aggregate consumption will in turn be given by C:

$$\begin{aligned} C &= \sum_{r=0}^{\infty}p^h\left(1-p^h\right)^r c\left(r\right) \qquad (4.27)\\ &= c\left(\infty\right)p^h\sum_{r=0}^{\infty}\left(1+\beta^r\overline{\mu}\right)^{\frac{1}{\sigma}} \end{aligned}$$

This must in turn be equal to the aggregate stock of resources available when assets are held constant at B:

$$C=p^hy^h+\left(1-p^h\right)y^l+\left(\beta^{-1}-1\right)B \tag{4.28}$$

Finally, if $\overline{\mu}$ is positive the participation constraint for high-income earners must be binding. Denoting their utility level U^h, if participation is optimal we must have:

$$U^h = \frac{c(0)^{1-\sigma} - 1}{1-\sigma} + \beta \left(p^h U^h + \left(1 - p^h\right) \left(\frac{c(1)^{1-\sigma} - 1}{1-\sigma} + \beta (...) \right) \right) \tag{4.29}$$

This implies:

$$U^h = \frac{1 - \beta \left(1 - p^h\right)}{1 - \beta} \sum_{r=0}^{\infty} \beta^r \left(1 - p^h\right)^r \frac{c(r)^{1-\sigma} - 1}{1-\sigma} \tag{4.30}$$

And the participation constraint itself gives:

$$U^h = \frac{\left(y^h\right)^{1-\sigma} - 1}{1-\sigma} + \frac{\beta}{1-\beta} \left[p^h \frac{\left(y^h\right)^{1-\sigma} - 1}{1-\sigma} + \left(1 - p^h\right) \frac{\left(y^l\right)^{1-\sigma} - 1}{1-\sigma} \right] \tag{4.31}$$

Together these conditions allow us to solve for $\overline{\mu}$ in terms of exogenous parameters, with solutions for the $c(r)$ consumption levels then following. We are able to confirm *ex post* that participation constraints are indeed satisfied for all agents. Clearly our ability to exploit such regularity in the evolution of the multipliers is dependent somewhat upon assuming just two possible shock realisations, but in principle the basic method can be generalised both to non-iid shocks and a larger set of income draws without too much difficulty. Moreover, the technical benefits of being able to focus exclusively on 'stochastic steady state' (which obtains, at least in this model, from the start of time whenever $R = \beta^{-1}$) are substantial: we have no reason to track the evolution of *non-stationary* multipliers through time, working instead simply with a distribution of these multipliers that is known to be time-invariant.

For any given y^h, y^l and p^h, high enough values of β will always ensure an equilibrium in which consumption is equalised across all agents and the participation constraints do not bind. This is because sufficiently patient agents will always prefer complete insurance in steady state, given their risk aversion – and when they are patient enough no feasible quantity of immediate inducements can be sufficient to offset the cost of losing this insurance in the longer term. By contrast, a combination of moderate impatience together with a value for y^h that is substantially greater than y^l but has a

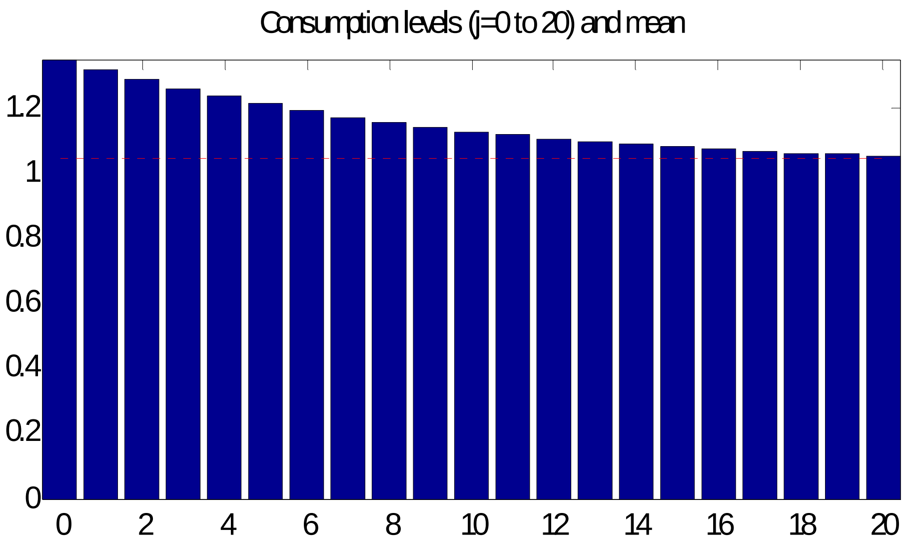

Figure 4.1: Consumption levels against periods since last high-income shock

low probability of occurring will ensure participation constraints bind. We simulate the model using illustrative values $y^h = 5$, $y^l = 1$, $p^h = 0.01$, $\beta = 0.9$, $\sigma = 1$, and assume the policymaker has no initial assets – implying an average per-capita consumption level of around 1.04 units.

4.4.2 Results

Figure 4.1 charts $c(r)$ for $r \in \{0, ..., 20\}$, together with a line denoting mean consumption. As can be seen, the policymaker ensures agents whose income is currently y^h enjoy a consumption level that is roughly 30 per cent above the mean value, but decays reasonably quickly through time (so long as subsequent income remains low) – consistent with the stationary character of the multipliers discussed above. The lowest consumption level (received by those whose have experienced incomes of y^l for a large number of consecutive periods) is still roughly 97 per cent of the mean. Figure 4.2 plots the Lorenz curve associated with this consumption distribution – evidencing considerable equality, with an associated Gini coefficient of 0.025.

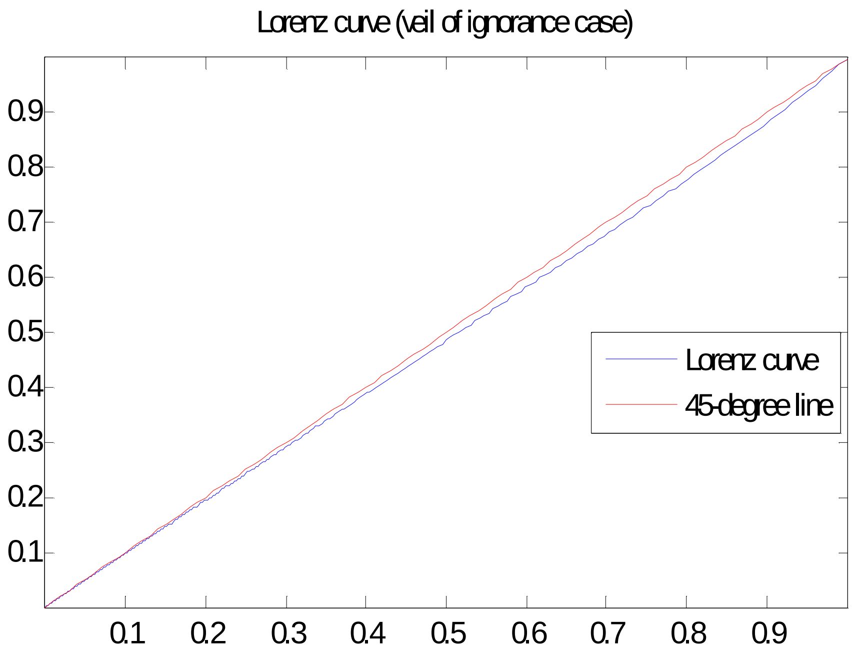

Figure 4.2: Lorenz curve under veil of ignorance policy

It is instructive to compare this outcome with a situation in which the time-0 optimal policy is implemented. The character of the solution in that case is well known (see Ljungqvist and Sargent (2004) for a textbook treatment): provided $R = \beta^{-1}$ continues to hold agents who receive an income of y^h in any time period from 0 onwards are forever rewarded with a *constant* consumption level that is just sufficient to incentivise their continued participation when they receive y^h. All remaining agents – constituting an ever-diminishing share of the total population – receive a lower consumption level, which is time-invariant. At the limit as time progresses every agent enjoys the high consumption level – a situation made possible by the policymaker having accumulated a sufficient quantity of assets in previous periods, which in turn is guaranteed by setting to a sufficiently low level the consumption levels of the agents yet to receive y^h.

We take the same calibration as for the veil of ignorance case, and again assume the policymaker's initial assets are zero. This implies a 'high' consumption level of 1.19 units – around 15 per cent above the mean level from the veil of ignorance case (note that in the case at hand mean consumption is increasing through time) – and a 'low' consumption level of 1.02 units – roughly 98.5 per cent of the prior mean. The distribution of consumption is changing through time: inequality by the Gini measure initially rises, as a select few agents gain access to permanently higher consumption entitlements, but then gradually falls as a diminishing number remain in the low consumption bracket.[82] The Gini coefficients associated with periods 0, 20, 50 and 100 are 0.002, 0.024, 0.037 and 0.034 respectively – so it takes a substantial time for the ultimate equality of outcomes to assert itself.

In the long run the policymaker must earn net interest on asset holdings each period equal to 15 per cent of the aggregate per-period endowment, in order to cover payments to those who have received a high income shock – implying a fairly sizeable long-run asset-to-output ratio of roughly 1.3. The time-inconsistency of the optimal plan derives from the incentive the policymaker has at dates subsequent to t to repudiate past promises and only give 'high' incomes to those receiving a

[82] One should remember that *aggregate* consumption is increasing at all horizons.

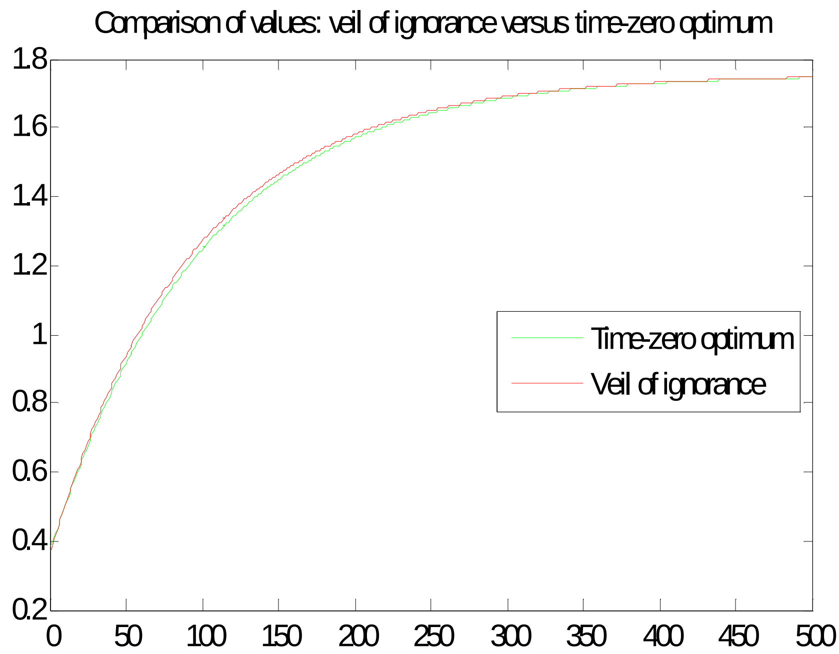

Figure 4.3: Relative values of veil of ignorance and time-zero policy

shock of y^h in the current time period – with the levels of consumption given to high- and low-income earners then capable of being redefined, reflecting the accumulation of assets since period 0. This would result in a redistribution towards those who previously would have had relatively little (since the 'low' consumption level will be higher now that assets have accumulated), which our utilitarian policymaker finds preferable. Since the optimal policy to implement at time t for a given quantity of assets B thereby depends upon the number of time periods since optimisation, agents at more distant horizons are being made subject to the particular perspective of time 0 – a problem that the veil of ignorance method avoids by design.

Figure 4.3 highlights this time inconsistency, charting the evolution of the values to the policymaker of a permanent switch to the veil of ignorance outcome and of continuing with the time-0 optimal outcome in perpetuity, with the former assessed *conditional upon the quantity of assets accumulated in*

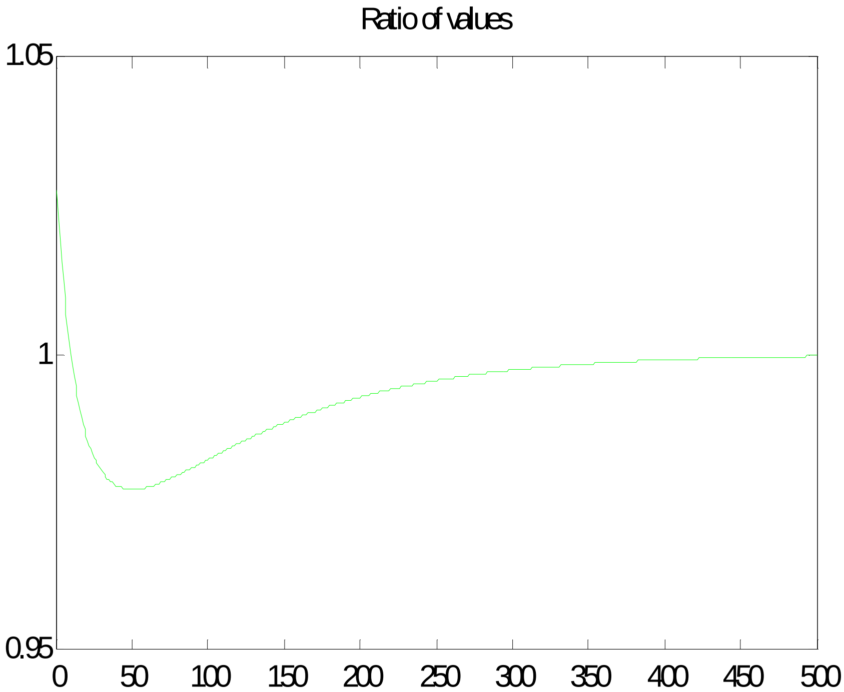

Figure 4.4: Ratio of values: veil of ignorance to time-zero policy

*each period along the time-*0 *path.*[83] Figure 4.4 plots the ratio of time-0 value to veil-of-ignorance value, assessed looking forward from each of the first 500 periods. Unsurprisingly, the time-0 value is initially the higher one (this must be true by definition), but not by all that much – and the ratio falls below 1 by period 9, where it forever remains. As time progresses the two values converge, simply because the quantity of assets accumulated under the time-0 optimal plan approaches the point where almost all agents are able to earn the same (high) consumption level without violating any participation constraints – an outcome that is first-best under both of the optimisation perspectives. The ratio reaches its lowest point around the 50th time period, when the time-0 plan is essentially dividing the population into two 'classes', based upon past endowment histories, and the merits of redistributing from those who have been lucky in the past to those who have not (which, to a greater extent than the time-0 plan, the veil of ignorance strategy does) are at their highest.

The comparison here should be treated with care: it is *not* the case that the veil of ignorance plan that follows from *zero* initial assets comes to dominate the time-0 plan at any horizon: the 'overtaking' exhibited in Figure 4.3 occurs only because we allow the veil of ignorance policymaker access to assets that he or she would not in fact have accumulated given the choice. But it does highlight how adherence to the time-0 plan may increasingly be seen as an impediment to more desirable, equitable outcomes as time progresses – and with scant intrinsic ethical justification for its policy evaluation perspective, it is perhaps unlikely to endure.

4.4.3 Comparison with alternative solution approaches

At this point we can again see the difficulty in generalising a 'timeless perspective' on optimality to this sort of setting. If we were to try to implement from time 0 the policy that would have been best from the perspective of a time period in the infinitely distant past we would face something of a conundrum:

[83] So, for instance, the time-0 optimum implies the policymaker carries around 0.25 units of assets into the 20th time period, and we can calculate the veil of ignorance policy associated with this initial asset stock – a policy that will then hold assets constant at 0.25 units in perpetuity.

the measure of agents previously having received a positive income shock between periods $-T$ and 0 clearly approaches unity as T grows (without bound) – implying that the timeless perspective ought to equalise incomes across agents *from the start of time*, whilst simultaneously allowing these incomes to be sufficiently high that incentive compatibility is satisfied for those with high current income draws.

But the policymaker at time 0 has yet to build up the stock of assets that would be needed to support in perpetuity a constant, universal consumption level sufficiently high to satisfy incentive compatibility for those currently receiving y^h. Thus any timeless perspective policymaker – indeed, any policymaker of any stripe – would be forced to vary consumption across agents in order to remain solvent (and keep those with high incomes from departing). But this variation could never have been optimal from the perspective of the distant past.

In a sense the problem is quite a fundamental one with the timeless perspective's normative perspective: what should one do if the best advice a policymaker from the distant past could give is: 'Don't start from here.'? A veil of ignorance approach avoids this problem – specifying an appropriate policy for any initial value of state variables (the asset stock in this case), irrespective of the decisions that may or may not have been taken by past policymakers applying the same methods.

Another perspective on this problem could potentially be obtained by applying the 'unconditionally optimal' method. This involves maximising the steady-state value of the policymaker's within-period welfare objective. But in the case at hand this turns out not to be possible. To see why, recall that the unconditionally optimal method was equivalent to a pursuing a time-0-optimal strategy under the assumption that the policymaker does not discount for pure time preference. This means that the first-order condition with respect to bond holdings *at any optimal steady state* when applying that method would have to be:

$$-\lambda_R(t) + R\lambda_R(t+1) = 0 \tag{4.32}$$

(where we again invoke a law of large numbers to eliminate aggregate uncertainty from the notation).

Since $R = \beta^{-1} > 1$, this equation implies $\lambda_R(t) > \lambda_R(t+1)$. That is, the shadow value of additional assets is declining at the optimum. But this in turn implies assets are still being accumulated – a result that is inconsistent with us being at an optimal steady state in the first place.

The basic problem here is that there simply doesn't exist a maximum possible value for the policymaker's within-period objective. Starting from any conjectured optimum it would always be possible to defer some consumption until later, and increase still further the value the objective in all future periods. This problem is likely to arise in a great deal of models: it depends only on the feasible asset stock for the policymaker being unbounded, which is likely to be true even in production economies if permanent growth is a possibility. Again, the veil of ignorance method retains the policymaker's basic time preference structure, so still considers deferred welfare of diminished importance relative to current welfare. And in any event, it is a method that remains well defined whether or not a stochastic steady state is obtained.

5 Example: A dynamic optimal income tax model

Part II of this thesis develops a novel approach for solving dynamic optimal tax models, of the kind popularised by the 'New Dynamic Public Finance' literature. It does so under the assumption that an optimal policy from the perspective of period 0 can always be implemented, despite the time-inconsistency of this solution. This time-inconsistency arises from the incentive of the policymaker in period 0 to promise high future utility to those who reveal themselves to have a relatively high type in the present – inducing a utility differential that it will not be desirable to maintain *ex post*.[84]

Moreover the time-inconsistency problem in this model is a particularly acute one, since under many preference specifications the long-run outcome is immiseration for almost all agents (as shown in Part II) – in the sense that consumption will converge almost surely to zero along almost all drawings for the dynamic productivity process, provided the real interest rate is equal to β^{-1}. There could be

[84] As studied in Part II, an additional source of time-inconsistency will result from the policymaker's incentive to reduce prior marginal benefits to mimicking in the event that productivity draws are not independent through time.

no starker illustration of a 'tyranny of the past' than this. Yet the immiseration result rests upon the central dynamic optimality condition in these models: the 'inverse Euler equation' – or its appropriate generalisation when preferences between consumption and labour supply are non-separable. If this equation does not apply under a 'veil of ignorance' optimality approach the immiseration results need not go through either. Even aside from normative questions, this would certainly support our argument that the veil of ignorance solution is likely to be more durable than the time-0 optimum.

5.1 Model assumptions

The basic model setup is presented in Part II of this thesis, so we choose not to repeat it here. Again we assume for simplicity that there exists a discrete, finite number N of possible within-period productivity types in the given set Θ, and will work under the assumption that the 'first-order' approach to the constraint set is valid in the current setting as there – so that conditional upon a given prior productivity history it is sufficient to restrict an agent of current type θ' to receive at least as much utility as he or she could obtain by mimicking an agent of type $\theta'' = \max\left\{\theta \in \Theta : \theta < \theta'\right\}$. We also assume that a continuum of agents exists (indexed on the interval $[0, 1]$), with independent stochastic (Markov) productivity processes across these agents. This implies the aggregate distribution of productivities across agents will almost surely be known in each period.

There is one important change that we must make to the modelling approach, and that is that individuals are characterised in all periods by a distinct drawing from the *infinite* set of past productivities, denoted $h_\theta \in \Theta^\infty$, and that this drawing for each agent is known to the policymaker. Since productivities are the relevant exogenous stochastic variables in this model, conditioning upon their infinite history will be desirable under the veil of ignorance solution. But it is clearly questionable whether this information will exist at time 0 – and the policymaker has no means to elicit it if it does not. Our assumption is effectively that agents' behaviour under whatever tax scheme has been in operation in the past can be used to infer their past types at all horizons, or at least to a sufficiently distant horizon that policy may very closely approximate the veil of ignorance optimum.[85] Given the detailed income tax records existing for individuals in many advanced economies this may not be so implausible as it first seems.

We interpret h_θ as the history of productivities for a given agent *not including their present draw.* Assuming perfect revelation in the past, this allows it to be associated with an agent's 'public' shock history, with current draws implicitly remaining private knowledge until the end of the time period. This provides a better mapping to our framework than if h_θ were to include present shocks, since incentive compatibility restrictions place cross restrictions on behaviour in response to *different* values of the current draw – a possibility not immediately admitted by the general forward-looking constraints of the form (2.2).

The complete history of exogenous variables h_ε then consists of a value for $h_\theta^i \in \Theta^\infty$ for every individual $i \in [0, 1]$, together with the history of real interest rates. The job of the veil of ignorance policymaker is then to choose consumption and output levels for each agent, together with a quantity of bond purchases, for each realisation of this h_ε and each quantity of inherited assets. As in the model of redistribution subject to participation constraints, it is simplest to assume that the real interest rate R is fixed at β^{-1} in all periods. By familiar logic, this together with the 'continuum of types' assumption ensures that each agent's allocation at a given point in time will depend only on his or her own type history, plus the aggregate stock of assets inherited from the previous time period; there is no aggregate uncertainty. We additionally assume that the exogenous unconditional probability measure over histories, π_Θ, is time-invariant – and thus on aggregate the policymaker faces the *same* distribution of productivity histories in each period.

[85] If we were to work with an overlapping generations model that featured no bequest motive we would only *need* knowledge on the (finite) prior productivity history of each agent alive at time 0.

5.2 Incorporating incentive compatibility constraints

In general we will suppress dependence on assets where appropriate, to keep the notation compact – though we introduce it in presenting the forward-looking policymaker's problem formally in the appendix. In writing incentive compatibility constraints it is also simplest to suppress the dependence of utility on consumption and output, linking it to actual and reported types alone. We then denote by $u(h_\theta, \theta)$ the within-period utility of an agent whose past report history (excluding current type) was h_θ, whose current type is θ, and who is truthfully reporting this (current) type. We denote by $\widehat{u}(h_\theta; \theta)$ the within-period utility of an agent with an identical set of past and present productivity draws, but who chooses in the present period to report $\theta' = \max\left\{\widetilde{\theta} \in \Theta : \widetilde{\theta} < \theta\right\}$ rather than θ. We then define truth-telling and mimicking value functions, $V(h_\theta; \theta)$ and $\widehat{V}(h_\theta; \theta)$ respectively, by the following recursions:

$$V(h_\theta; \theta) = u(h_\theta; \theta) + \beta \sum_{\theta'' \in \Theta} \pi_\Theta(\theta'' | \theta) V((h_\theta, \theta); \theta'') \tag{5.1}$$

$$\widehat{V}(h_\theta; \theta) = \widehat{u}(h_\theta; \theta) + \beta \sum_{\theta'' \in \Theta} \pi_\Theta(\theta'' | \theta) V((h_\theta, \theta'); \theta'') \tag{5.2}$$

where again $\theta' = \max\left\{\widetilde{\theta} \in \Theta : \widetilde{\theta} < \theta\right\}$, and the argument (h_θ, θ) used in $V((h_\theta, \theta); \theta'')$ is the infinite history given by h_θ together with an additional final entry of θ. We appeal to the Markov assumption in the second of these expressions to assert that the value to an agent of being type θ'' with a report history (h_θ, θ) is independent of whether that history was *truthfully* reported. Note that $\widehat{V}(h_\theta, \theta)$ is not defined when θ is the minimal element of Θ, which we denote $\underline{\theta}$ throughout. (We similarly denote by $\overline{\theta}$ the maximal element of Θ.)

The 'downwards' incentive compatibility restriction that we assume to be binding is then, for each

$(h_\theta;\theta) \in \Theta^\infty$ where θ is not the minimal element of Θ:

$$V(h_\theta;\theta) \geq \widehat{V}(h_\theta;\theta) \tag{5.3}$$

to which is added the 'forward-looking' recursive definitions of $V(h_\theta;\theta)$ and $\widehat{V}(h_\theta;\theta)$ in characterising the set of forward-looking constraints. These definitions are not yet exactly consistent with the general format of (2.2), since they do not contain the equivalent of g_1 functions for *all* successor histories to h_θ. When writing the policymakers' objectives one can instead substitute the following:

$$\begin{aligned} & V(h_\theta;\theta)\,\pi_\Theta(\theta|h_\theta) - u(h_\theta;\theta)\,\pi_\Theta(\theta|h_\theta) \\ = & \sum_{\widetilde{\theta}\in\Theta} \pi_\Theta\left(\widetilde{\theta}|h_\theta\right)\beta \sum_{\theta''\in\Theta} I\left(\widetilde{\theta}=\theta\right)\pi_\Theta(\theta''|\theta)\,V((h_\theta,\theta);\theta'') \end{aligned} \tag{5.4}$$

and:

$$\begin{aligned} & \widehat{V}(h_\theta;\theta)\,\pi_\Theta(\theta'|h_\theta) - \widehat{u}(h_\theta;\theta)\,\pi_\Theta(\theta'|h_\theta) \\ = & \sum_{\widetilde{\theta}\in\Theta} \pi_\Theta\left(\widetilde{\theta}|h_\theta\right)\beta \sum_{\theta''\in\Theta} I\left(\widetilde{\theta}=\theta'\right)\pi_\Theta(\theta''|\theta)\,V((h_\theta,\theta');\theta'') \end{aligned} \tag{5.5}$$

where $I\left(\widetilde{\theta}=\theta\right)$ is an indicator function taking the value 1 when $\widetilde{\theta}=\theta$ and 0 otherwise, and θ' is defined as before. We have then worked the state transition probabilities $\pi_\Theta\left(\widetilde{\theta}|h_\theta\right)$ fully into the constraints, as demanded by (2.2). The first line of each of the preceeding equations then corresponds to a function of the form g_0, and the associated equivalent of the g_1 function will be (minus) the terms in the second line from β onwards. For each history h_θ there are then N constraints of the form (5.4) (one for each $\theta \in \Theta$), and $N-1$ of the form (5.5).(one for each $\theta \in \Theta \backslash \underline{\theta}$).

5.3 Optimal policy: mean reversion in place of immiseration

The outcome of the backward-looking policymaker's problem is relatively straightforward, and we neglect to write out the relevant objective in full. Choice is only over the quantity of assets, B, to hold in each period, and since $R = \beta^{-1}$ and there is no aggregate uncertainty, optimal decisions can be expressed by a first-order condition identical to equation (4.23) that we obtained in the model of redistribution under imperfect commitment:

$$-\lambda_R (B) + \lambda_R (B (B)) = 0 \tag{5.6}$$

where λ_R is the multiplier on the within-period resource constraint, and $B(\cdot)$ is the policy function for the optimal choice of bonds. Since the shadow value to the *forward*-looking policymaker of additional resources will be diminishing in assets, this equation implies that at any veil of ignorance optimum wealth will be held constant: $B(B) = B$.

Optimal forward-looking policy requires more detailed consideration, with some fairly involved (though not difficult) algebra. We relegate the main arguments to an appendix, where we show that the following condition must hold at a veil of ignorance optimum, for all h_θ and θ:

$$\begin{aligned} & \sum_{\theta' \in \Theta} \pi_\Theta (\theta' | h_\theta) \left[MC \left((h_\theta, \theta); \theta' \right) - \overline{MC} \right] \\ = \; & \beta \left[MC (h_\theta; \theta) - \overline{MC} \right] \end{aligned} \tag{5.7}$$

where we define $MC(h_\theta; \theta)$ as the marginal cost of incentive-compatible utility provision to an agent of current type θ and history h_θ, familiar from Part II:

$$MC (h_\theta; \theta) \equiv \frac{1 - \alpha (h_\theta; \theta)}{u_c (h_\theta; \theta) + u_y (h_\theta; \theta) \alpha (h_\theta; \theta)} \tag{5.8}$$

where the relevant definition of $\alpha(h_\theta;\theta)$ in this case is:

$$\alpha(h_\theta;\theta) \equiv \frac{u_c(h_\theta;\theta) - \widehat{u}_c(h_\theta;\theta')}{\widehat{u}_y(h_\theta;\theta') - u_y(h_\theta;\theta)} \tag{5.9}$$

for $\theta' = \min\left\{\widetilde{\theta} \in \Theta : \widetilde{\theta} > \theta\right\}$. (If $\theta = \max\left\{\widetilde{\theta} \in \Theta\right\}$ we simply define $\alpha(h_\theta;\theta) = 0$.) $\overline{MC}$ is then the population average of $MC(h_\theta;\theta)$, taken across all past and current productivity draws, using the measure π_Θ. In general higher values for this object will imply that agents are better off (it is more costly on average to provide them with utility at the margin), and thus will be associated with a higher initial asset stock.

Condition (5.7) is just one of the optimality conditions that will characterise the veil of ignorance solution to this model, but it is exceptionally insightful on the long-run properties of that solution. Recall from Part II that a time-0 optimum when $R = \beta^{-1}$ satisfies a martingale condition of the form:

$$MC_t(h_\theta;\theta) = \sum_{\theta' \in \Theta} \pi_\Theta(\theta'|h_\theta)\, MC_{t+1}((h_\theta,\theta);\theta') \tag{5.10}$$

It was from this result that we were able to prove 'immiseration' results, by applying Doob's martingale convergence theorem whenever a zero lower bound could be put on the marginal cost of incentive-compatible utility provision.

Instead of a martingale, the veil of ignorance optimum requires the marginal cost of utility provision to be *mean-reverting* for each agent through time: according to (5.7) the expected value of the marginal cost across productivity draws for agents with a common prior history will deviate from its population mean value by a fration β of its deviation in the previous period. And since the distribution of histories is not changing through time, the distribution of pre-tax income and consumption in the population must likewise remain stable through time, with no long-run tendency towards immiseration for any positive measure of agents.

Notice the strong parallels between the outcomes in this model and the analysis of redistribution under participation constraints. There too, when the interest rate was equal to the inverse discount factor the veil of ignorance optimum involved mean-reverting outcomes for each member of the population, with *aggregate* wealth held constant. By contrast, in both cases the time-0 optimum involved non-stationary outcomes at the individual level – along with changes to the aggregate asset stock through time.[86]

Why does equation (5.7) hold? The basic intuition is best seen by recalling how (5.10) was derived. The idea there was that we could always increase utility at t for agents with a given history (and a given current productivity draw), and reduce it uniformly across all of the agents at $t+1$ who experienced that same history up to t. If the latter utility reduction was equal (in discounted present value) to the earlier increase, and if the perturbations took place along the appropriate dimensions in consumption-output space, incentive compatibility would be preserved – at least for the relaxed problem. Since the (utilitarian) policymaker's objective would be left unchanged by these changes, we could only be at an optimum if the marginal cost of utility provision at t for agents with a given history were equal to its expected value across types at $t+1$ (allowing for the effects of interest and time preference discounting – which cancel when $R = \beta^{-1}$).

When we consider a veil of ignorance optimum similar logic can be applied, at least informally. In principle one could always increase by a unit the within-period utility given to agents with a particular (infinite) type history h_θ and current type θ, and reduce the within-period utility of all agents with *past* histories of (h_θ, θ) by β^{-1} units. With movements along appropriate dimensions in consumption-output space (changing the utility of 'downwards mimickers' by an identical amount to truth-tellers), this perturbation could preserve incentive compatibility (under the relaxed problem) for all concerned. But a veil of ignorance policymaker must behave as though the *successor* histories to (h_θ, θ) are equally likely to be observed in each period as is (h_θ, θ) itself. Hence the perturbation will be considered to reduce *aggregate* utility by an amount $\beta^{-1} - 1$ (per agent with history (h_θ, θ)) – since the β^{-1} utility

[86] The latter is not immediate from (5.10), but aggregate assets do indeed tend to decline through time in the computed example of Part II.

reduction occurs *simultaneously* to the unit increase from the policymaker's perspective. This can only be consistent with an optimal allocation if there is an offsetting net marginal increase in the policymaker's *resources* (assessed in utility units) as a consequence of the two perturbations. From these considerations we can write the condition:

$$\begin{aligned} & 1-\beta^{-1} \\ = & \frac{1}{\overline{MC}}\left[MC\left(h_{\theta};\theta\right)-\beta^{-1}\sum_{\theta'\in\Theta}\pi_{\Theta}\left(\theta'|h_{\theta}\right)MC\left(\left(h_{\theta},\theta\right);\theta'\right)\right] \end{aligned} \tag{5.11}$$

The left-hand side here is the net effect of the two perturbations on utility (per agent with with the relevant history), and the right-hand side is their net marginal cost, assuming that these can be converted into units of utility by dividing by the population average marginal cost of utility provision, $\overline{MC}$.[87] A simple rearrangement of this condition would indeed give (5.7). Though the full derivation of (5.7) in the appendix is considerably more involved than this, the same basic intuition lies behind it.

Condition (5.7) is instructive on both the normative and positive benefits of a veil of ignorance approach. Since we need no longer expect long-run immiseration, the 'tyranny of the past' implicit in that outcome has been avoided, without sacrificing entirely the benefits from spreading incentive structures throught time. This is exactly in keeping with the basic Rawlsian perspective, that opportunities for social improvement should be exploited wherever possible when designing social institutions – but only to the extent that they would be identified as improvements even by agents who disregarded the particularities of their individual circumstances (in this case the time period in which they choose). From a positive perspective, the likelihood of long-run endurance for a tax system that does *not* imply complete immiseration for almost all agents in the long run is surely greater. The veil of ignorance

[87] Uniform utility provision to all agents in the population simultaneously along an appropriate dimension in consumption-output space would preserve incentive compatibility. $\overline{MC}$ will be the average cost of this at the margin. Hence $\frac{1}{\overline{MC}}$ is the marginal utility value to the policymaker of a unit of resources.

policy seems much more likely to be *implementable*.

As in the redistribution model, alternative 'time-invariant' solution methods are not well defined here. Since the real interest rate is fixed at β^{-1}, any 'unconditionally optimal' solution generally encounters the problem that an *optimal* steady state does not exist. Average welfare will generally be increasing in assets, whilst any steady state must be holding assets constant. There is no upper bound on the feasible asset stock, so there will be no maximum constant asset level – and thus no unconditional optimum.

As for the timeless perspective, presumably this would imply immiseration for (almost) all agents from the *start* of time. Quite aside from the fact that it is clearly undesirable, this allocation could additionally imply a dynamic resource surplus for the policymaker (a counterpart to resource *deficiency* that prevented timeless perspective policy from working in the redistribution model): Part II showed that 'no distortion at the top' holds at all horizons under the time-0 optimum, and thus a positive measure of agents will certainly be producing output under an immiserating allocation. Unless the long-run consumption of a zero-measure subset of fortunate agents is very high indeed, the reason for long-run immiseration under the time-0 optimum is that the policymaker chooses to run down assets, and the limiting allocation is a product of the escalating requirements of debt servicing. The timeless perspective policymaker would then be servicing debt without owing it – a wasteful endeavour, to put it mildly.

5.4 Other aspects to optimal policy

Given that the veil of ignorance approach chiefly affects the treatment of dynamics, it is perhaps unsurprising that the 'intratemporal' optimality requirements derived in Part II of this thesis carry over in their entirety. Suppose that for a given history h_θ we have an N-dimensional 'utility increment' vector ν, with each entry corresponding to one of the elements of Θ, satisfying:

$$\sum_{\theta\in\Theta}\pi_{\Theta}\left(\theta|h_{\theta}\right)\nu\left(\theta\right)=\sum_{\theta\in\Theta}\pi_{\Theta}\left(\theta|h'_{\theta}\right)\nu\left(\theta\right)=0 \tag{5.12}$$

where h'_{θ} is the history obtained by replacing the last entry in h_{θ} with the element one higher in Θ (if the last entry in h_{θ} is the maximal element of Θ then let $h_{\theta}=h'_{\theta}$). Then a necessary condition for the forward-looking policymaker to be behaving optimally is the following:

$$\begin{aligned}&\sum_{\theta\in\Theta}\pi_{\Theta}\left(\theta|h_{\theta}\right)\nu\left(\theta\right)MC\left(h_{\theta};\theta\right)\\=&\sum_{\theta\in\Theta\backslash\overline{\theta}}\pi_{\Theta}\left(\theta|h_{\theta}\right)\left(\nu\left(\theta'\right)-\nu\left(\theta\right)\right)TC\left(h_{\theta};\theta\right)\end{aligned} \tag{5.13}$$

where the term $\nu\left(\theta'\right)$ in the last sum is the entry in ν corresponding to a productivity draw one higher than the given θ, and $TC\left(h_{\theta};\theta\right)$ is a marginal resource cost associated with the within-period income tax wedge faced by the agent of type $\left(h_{\theta};\theta\right)$, defined by:

$$TC\left(h_{\theta};\theta\right)\equiv\frac{\tau\left(h_{\theta};\theta\right)}{\widehat{u}_{c}\left(h_{\theta};\theta'\right)\left(1-\tau\left(h_{\theta};\theta\right)\right)+\widehat{u}_{y}\left(h_{\theta};\theta'\right)} \tag{5.14}$$

for $\theta'=\min\left\{\widetilde{\theta}\in\Theta:\widetilde{\theta}>\theta\right\}$ and $\tau\left(h_{\theta};\theta\right)=1+\frac{u_{y}(h_{\theta};\theta)}{u_{c}(h_{\theta};\theta)}$. As discussed in Part II, this is the marginal cost to the policymaker of a movement 'down' the within-period indifference curve of an agent of type $\left(h_{\theta};\theta\right)$ by an amount just sufficient to reduce the welfare of mimickers by a unit.

A derivation of (5.13) is provided in the appendix. Though the exact presentation differs (we have not used matrix algebra here), it is exactly equivalent to equation (6.4), derived in Part II of this thesis. The reader is referred there for a full discussion. When shocks are iid or the last entry of h_{θ} is $\overline{\theta}$, condition (5.13) provides $N-1$ restrictions (one for each possible linearly independent ν vector); otherwise it gives $N-2$ conditions. The only major difference here from the necessary conditions for time-0 optimality is that we will not have an initial period in which there are $N-1$ restrictions provided by (5.13) for the more general Markov case with $\theta \neq \overline{\theta}$.

It is also straightforward to show that 'no distortion at the top' still applies, and so $\tau\left(h_{\theta};\overline{\theta}\right)=0$. In general this still leaves us one restriction short of tying down allocations fully. As before, the additional equation will come from considering the most efficient way to reduce the 'information rents' that higher types can earn – exploiting differences in the probability measures across future states applied by mimickers and truth-tellers. This yields the following optimality condition (again derived in the appendix), which must hold for all $\theta \neq \overline{\theta}$ and all $h_{\theta} \in \Theta^{\infty}$:

$$\begin{aligned}\beta TC\left(h_{\theta};\theta\right) &= \sum_{\theta''\in\Theta\backslash\overline{\theta}} \pi_{\Theta}\left(\theta''\right|\left(h_{\theta},\theta\right)\right)\left(\nu\left(\theta'\right)-\nu\left(\theta''\right)\right)TC\left(\left(h_{\theta},\theta\right);\theta''\right) \\ &\quad -\sum_{\theta''\in\Theta} \pi_{\Theta}\left(\theta''\right|\left(h_{\theta},\theta\right)\right)\nu\left(\theta''\right)MC\left(\left(h_{\theta},\theta\right);\theta''\right)\end{aligned} \tag{5.15}$$

(where $\nu\left(\theta'\right)$ is now taken to be the entry in ν one higher than $\nu\left(\theta''\right)$), for any ν vector such that

$$\sum_{\theta''\in\Theta} \pi_{\Theta}\left(\theta''\right|\left(h_{\theta},\theta\right)\right)\nu\left(\theta''\right)=0$$

and

$$\sum_{\theta''\in\Theta} \pi_{\Theta}\left(\theta''\right|\left(h_{\theta},\theta^{+}\right)\right)\nu\left(\theta''\right)=1$$

where $\theta^{+}=\min\left\{\widetilde{\theta}\in\Theta:\widetilde{\theta}>\theta\right\}$.

This is the natural analogue to expression (6.5) derived in Part II. Once again, if the policymaker is distorting the production of type $(h_{\theta};\theta)$ by a significant amount, this must be because there are significant marginal benefits to reducing the income that must be paid to higher types. This 'information rent' can be reduced still further if the allocation of utility across successor histories to $(h_{\theta};\theta)$ can be skewed in a manner that benefits prior truth-tellers more than prior mimickers. The greater is the distortion $TC(h_\theta;\theta)$, the greater is this 'twist' to allocations under successor histories to $(h_\theta;\theta)$ – the marginal cost of which is given by the right-hand side of (5.15). The only significant difference between (5.15) and its time-0 analogue is the appearance of β pre-multiplying the object on the left-hand side. Under the time-0 optimum with $R = \beta^{-1}$ this term would disappear (more generally it would be replaced by $R\beta$ instead).

Once again, the veil of ignorance optimum allows for a diminution through time at rate β in the size of distortions introduced for the sake of prior incentive compatibility constaints. Indeed, equation (5.15) is perhaps most interesting in the context of the recent work by Farhi and Werning (2010), which shows substantial upward drift in the average labour distortion through time, under a time-0 optimum for an overlapping generations version of the dynamic Mirrlees setup. It seems likely that a condition similar to (5.15) will still characterise dynamic optimality in a veil of ignorance optimum for their setup, and the 'deadening' effect of the additional β would almost certainly reduce in the speed at which implicit marginal taxes rise, or even prevent this rise altogether.

6 Conclusion

This paper investigates the consequences for time-inconsistency problems of assuming that optimal policy must be chosen without any knowledge of the time period in which it is to be implemented. The purpose of doing so is to provide an alternative to standard 'commitment' solutions to these problems, which assume that the will of a policymaker who exists at the very start of time can be imposed on all of his or her successors regardless how imperfect the associated policies have come to appear. Whilst it may be plausible in some settings for legal contracts to permit 'once and for all' choice, for most policy questions this is not true, and there are grounds for exploring alternative approaches that will respect a need to abide by expectational constraints without appearing partial to one particular optimality perspective. We make particular appeal to the social contractarian tradition in political theory, and the idea (famously articulated by Rawls (1971)) that policymaking institutions should be designed without reference to the designers' particular, transitory circumstances. Since (in the models that we have considered) the time period in which a given policymaker exists fully characterises what is particular about his or her perspective, it is natural to ask what are the consequences setting policy without explicit knowledge of time.

The policies generated by this solution approach are best interpreted as the sorts of measures that a Rawlsian institutional designer would wish to task independent bodies (such as an independent central bank) with implementing, or would wish to incorporate into relevant taxation or social insurance legislation – the point being that the parsimony inherent in our method makes its prescriptions far more likely to survive interrogation by subsequent policymakers (and their publics) than legislation drafted in the specific best interest of one particular government.

From an economic perspective our approach exhibits an interesting asymmetry. We show that denying knowledge of time when choosing variables that feature in expectational constraints eliminates the conventional 'inverse discounting' of expectational effects. This is because any given history is just as likely to characterise the first time period or the second, and thus choices in response to that history are equally likely to have a direct effect on the underlying policy objective in the first period as they are to have an effect via first-period *expectations* of outcomes in the second period – which in turn affect possible first-period choices. By contrast, when choice is time-specific any one-period-ahead expectation channel will be known to have an effect on constraints the period *before* it has a direct effect on welfare, and thus the expectational effects are weighted more highly relative to the direct effects, due to pure time preference.

The asymmetry arises because there will still be an incentive in our setup to discount the effects on *future* policy possibilities of changing state variables that feature in 'backward-looking' constraints. It does not take *specific* knowledge of the time period to know that increases to the current value of the capital stock, for instance, will affect production possibilities one period after the sacrifices required for higher accumulation – and the time preference weightings applied by the policymaker to distinct marginal effects reflect this.

Thus our method achieves something of a ‘third way’ between policy approaches that maximise the expected value of steady state welfare – as if the policymaker never discounts the future – and those that apply standard discounting for pure time preference. Where ‘forward-looking’ variables feature we are closer to the former, and where ‘backward-looking’ variables feature we are closer to the latter.

Moreover, we show that our ‘veil of ignorance’ approach can be implemented across a far broader set of models than the main alternative (normative) proposal for overcoming time inconsistency in the literature, the ‘timeless perspective’ due to Woodford (2003). That method has particular difficulties when infinite-horizon utility constraints feature, whereas these are very straightforward to incorporate under our approach. There are also clear normative advantages to what we propose: we show that in one very simple class of models the timeless perspective policy is Pareto-dominated by the veil of ignorance policy, in the sense that policymakers in *all* time periods would prefer the latter to the former.

More contentiously, we argue that the veil of ignorance approach will often be more likely to endure future democratic choice than optimal plans from the perspective of an initial time period. These optimal plans are well known in some models to result in very extreme long-run outcomes, under conventional assumptions about the real rate of interest – a notable example being dynamic optimal taxation models, where long-run ‘immiseration’ has been shown in Part II of this thesis to occur (almost surely) for each agent under a wide class of preference assumptions. When the veil of ignorance approach to optimality is applied instead, allocations to each agent satisfy a simple mean-reversion property over time, making it far less likely that future generations will come to resent – and to overthrow – the regime imposed by their ancestors.

A Appendix to Part II

A.1 Proof of property 4 of first-best allocation (decreasing utility in type)

It is useful first to show that normality of leisure implies $u_{cc} + u_{cy} < 0$. The consumer's within-period problem (at autarky prices) is:

$$\max_{c,y} u(c, y; \theta)$$

subject to

$$c = y + \omega$$

for some endowment ω. At any interior optimum we will have $u_c = -u_y$, and differentiating both this and the budget constraint totally with respect to ω gives:

$$\frac{dy}{d\omega} = -\frac{u_{cc} + u_{cy}}{u_{cc} + 2u_{cy} + u_{yy}}$$

The denominator here is negative by the negative definiteness of the partial Hessian, and so $\frac{dy}{d\omega} < 0$ only if $u_{cc} + u_{cy} < 0$, as required.

We can now analyse the impact of an increase in θ on first-best outcomes by taking a total derivative of utility with respect to θ, under the twin restrictions (valid at the first-best) that u_c and u_y remain unchanged. With a little algebra it can be shown that these restrictions imply:

$$\frac{dc}{d\theta} = \frac{u_{yy}u_{c\theta} - u_{cy}u_{y\theta}}{u_{cy}^2 - u_{cc}u_{yy}}$$

$$\frac{dy}{d\theta} = \frac{u_{cc}u_{y\theta} - u_{cy}u_{c\theta}}{u_{cy}^2 - u_{cc}u_{yy}}$$

The overall effect on utility at the margin is:

$$\begin{aligned}\frac{du}{d\theta} &= u_\theta + u_c\frac{dc}{d\theta} + u_y\frac{dy}{d\theta} \\ &= u_\theta + u_y\frac{u_{y\theta}\left(u_{cc}+u_{cy}\right) - u_{c\theta}\left(u_{yy}+u_{cy}\right)}{u_{cy}^2 - u_{cc}u_{yy}}\end{aligned}$$

(where we have used that $u_c = -u_y$ at the optimum). Negative definiteness of the partial Hessian $\begin{bmatrix} u_{cc} & u_{cy} \\ u_{cy} & u_{yy}\end{bmatrix}$ implies $u_{cy}^2 - u_{cc}u_{yy} < 0$,[88] and we have $u_y < 0$, so if $u_{cc} + u_{cy} < 0$ then condition (2.2) gives:

$$\begin{aligned}\frac{du}{d\theta} &< u_\theta + u_y\frac{u_{yy}\frac{u_\theta}{u_y}\left(u_{cc}+u_{cy}\right) - u_{c\theta}\left(u_{yy}+u_{cy}\right)}{u_{cy}^2 - u_{cc}u_{yy}} \\ &= u_\theta\left(1 + \frac{u_{yy}\left(u_{cc}+u_{cy}\right) - u_{cy}\left(u_{yy}+u_{cy}\right)}{u_{cy}^2 - u_{cc}u_{yy}}\right) \\ &= 0\end{aligned}$$

where we have additionally used condition (2.3). The result then follows immediately.

A.2 Proof of Proposition 1

For the sake of clarity we index the N elements of Θ in ascending order, so $\theta_t^n > \theta_t^m$ whenever $n > m$ for all $n, m \in \{1, ..., N\}$. We have imposed that

$$W\left(\theta_t^n; \theta_t^n, \widehat{\theta}^{t-1}\right) = W\left(\theta_t^{n-1}; \theta_t^n, \widehat{\theta}^{t-1}\right)$$

for all $n \in \{2, ..., N\}$, and wish to show that this implies

$$W\left(\theta_t^n; \theta_t^n, \widehat{\theta}^{t-1}\right) \geq W\left(\theta_t^m; \theta_t^n, \widehat{\theta}^{t-1}\right)$$

[88] Consider movements according to the vector $\begin{bmatrix} -\frac{u_{yc}}{u_{cc}} \\ 1 \end{bmatrix}$ here.

for all $m \in \{1, ..., N\}$, given the increasing differences condition

We first consider the case in which $n > 1$, and show

$$W\left(\theta_t^n; \theta_t^n, \widehat{\theta}^{t-1}\right) \geq W\left(\theta_t^m; \theta_t^n, \widehat{\theta}^{t-1}\right)$$

for all $m \in \{1, ..., n-1\}$. For $m = n-1$ this holds by assumption. For $m = n-2$ we have by increasing differences:

$$\begin{aligned} & W\left(\theta_t^{n-1}; \theta_t^n, \widehat{\theta}^{t-1}\right) - W\left(\theta_t^{n-1}; \theta_t^{n-1}, \widehat{\theta}^{t-1}\right) \\ > \; & W\left(\theta_t^{n-2}; \theta_t^n, \widehat{\theta}^{t-1}\right) - W\left(\theta_t^{n-2}; \theta_t^{n-1}, \widehat{\theta}^{t-1}\right) \end{aligned}$$

But

$$W\left(\theta_t^{n-1}; \theta_t^n, \widehat{\theta}^{t-1}\right) = W\left(\theta_t^n; \theta_t^n, \widehat{\theta}^{t-1}\right)$$

and

$$W\left(\theta_t^{n-2}; \theta_t^{n-1}, \widehat{\theta}^{t-1}\right) = W\left(\theta_t^{n-1}; \theta_t^{n-1}, \widehat{\theta}^{t-1}\right)$$

so prior inequality implies

$$W\left(\theta_t^n; \theta_t^n, \widehat{\theta}^{t-1}\right) > W\left(\theta_t^{n-2}; \theta_t^n, \widehat{\theta}^{t-1}\right)$$

as required. Taking $m = n-3$, we then have by increasing differences:

$$\begin{aligned} & W\left(\theta_t^{n-2}; \theta_t^n, \widehat{\theta}^{t-1}\right) - W\left(\theta_t^{n-2}; \theta_t^{n-2}, \widehat{\theta}^{t-1}\right) \\ > \; & W\left(\theta_t^{n-3}; \theta_t^n, \widehat{\theta}^{t-1}\right) - W\left(\theta_t^{n-3}; \theta_t^{n-2}, \widehat{\theta}^{t-1}\right) \end{aligned}$$

Again, by

$$W\left(\theta_t^{n-3}; \theta_t^{n-2}, \widehat{\theta}^{t-1}\right) = W\left(\theta_t^{n-2}; \theta_t^{n-2}, \widehat{\theta}^{t-1}\right)$$

this inequality collapses to

$$W\left(\theta_t^{n-2};\theta_t^n,\widehat{\theta}^{t-1}\right) > W\left(\theta_t^{n-3};\theta_t^n,\widehat{\theta}^{t-1}\right)$$

and we can apply the earlier result

$$W\left(\theta_t^{n};\theta_t^n,\widehat{\theta}^{t-1}\right) > W\left(\theta_t^{n-2};\theta_t^n,\widehat{\theta}^{t-1}\right)$$

to assert

$$W\left(\theta_t^{n};\theta_t^n,\widehat{\theta}^{t-1}\right) > W\left(\theta_t^{n-3};\theta_t^n,\widehat{\theta}^{t-1}\right)$$

as required. The same argument can clearly be applied for all $m \in \{1,...,n-1\}$.

When $n < N$ we must in the same way consider the cases of $m \in \{n+1,...,N\}$. For $m = n+1$, we have immediately by the binding restriction on $n+1$-types, together with increasing differences:

$$\begin{aligned} 0 &= W\left(\theta_t^{n+1};\theta_t^{n+1},\widehat{\theta}^{t-1}\right) - W\left(\theta_t^{n};\theta_t^{n+1},\widehat{\theta}^{t-1}\right) \\ &> W\left(\theta_t^{n+1};\theta_t^{n},\widehat{\theta}^{t-1}\right) - W\left(\theta_t^{n};\theta_t^{n},\widehat{\theta}^{t-1}\right) \end{aligned}$$

as required. By similar logic, for $m = n+2$ we have:

$$\begin{aligned} 0 &= W\left(\theta_t^{n+2};\theta_t^{n+2},\widehat{\theta}^{t-1}\right) - W\left(\theta_t^{n+1};\theta_t^{n+2},\widehat{\theta}^{t-1}\right) \\ &> W\left(\theta_t^{n+2};\theta_t^{n},\widehat{\theta}^{t-1}\right) - W\left(\theta_t^{n+1};\theta_t^{n},\widehat{\theta}^{t-1}\right) \end{aligned}$$

and the condition

$$W\left(\theta_t^{n};\theta_t^n,\widehat{\theta}^{t-1}\right) > W\left(\theta_t^{n+1};\theta_t^n,\widehat{\theta}^{t-1}\right)$$

then delivers the required result. Again, we can apply an identical argument inductively for all remaining $m < N$. This completes the proof.

A.3 Proof of Proposition 6

We again consider a pair of perturbation schedules $\Delta_{-1}(\delta)$ and $\Delta(\delta)$ applied at t and $t+1$ respectively to the allocations of agents with the relevant reporting history θ^t. We set the first $n-1$ rows of $\Delta(\delta)$ to 0 for all δ. The nth row is then constructed to equal

$$\left(\varphi^c\left(\theta_{t+1}^n, \delta; c_{t+1}^*, y_{t+1}^*\right), \varphi^y\left(\theta_{t+1}^n, \delta; c_{t+1}^*, y_{t+1}^*\right)\right)$$

where these values are defined implicitly (and – by the single-crossing condition – uniquely) by the following two equations:

$$u\left(c+\varphi^c(\theta, k; c, y), y+\varphi^y(\theta, k; c, y); \theta\right) = u(c, y; \theta) \tag{A.1}$$

$$u\left(c+\varphi^c(\theta, k; c, y), y+\varphi^y(\theta, k; c, y); \theta'\right) = u\left(c, y; \theta'\right) + k \tag{A.2}$$

with $\theta' = \min\left\{\theta'' \in \Theta : \theta'' > \theta\right\}$. (So $\varphi^c(\theta, k; c, y)$ and $\varphi^y(\theta, k; c, y)$ are perturbations from the allocation (c, y) that keep utility constant for an agent of type θ, whilst increasing it by k units for an agent whose type is one higher.) The mth row in $\Delta(\delta)$ is given for $m \in \{n+1, ..., N\}$ by

$$\left(\phi^c\left(\theta_{t+1}^m, \delta; c_{t+1}^*, y_{t+1}^*\right), \phi^y\left(\theta_{t+1}^m, \delta; c_{t+1}^*, y_{t+1}^*\right)\right)$$

where these functions are defined in the proof of Proposition 3. We additionally assume a time-t perturbation schedule $\Delta_{-1}(\delta)$ given by:

$$\begin{aligned} & \Delta_{-1}(\delta) \\ = \quad & \left(\phi^c\left(\theta_t, -\beta\delta\pi_\Theta\left(\theta_{t+1} > \theta_{t+1}^n\right); c_t^*, y_t^*\right), \phi^y\left(\theta_t, -\beta\delta\pi_\Theta\left(\theta_{t+1} > \theta_{t+1}^n\right); c_t^*, y_t^*\right)\right) \end{aligned} \tag{A.3}$$

These perturbations will preserve incentive compatibility (according to the relaxed problem's constraint set) at $t+1$. For the nth agent this holds because the perturbation is constructed so as not to affect his or her utility from truth-telling, whilst utility from mimicking the $n-1$th agent is held constant by the fact agents below the nth see no change to their allocations. For agents whose types are higher than the nth, the $\Delta(\delta)$ schedule is constructed to ensure there are equal utility gains to mimicking and truth-telling, so that 'downwards' incentive compatibility constraints cannot be violated through these perturbations, whilst we are free to ignore other constraints under the supposition that the general problem's optimum cannot be improved upon by any allocation that satisfies the relaxed problem's constraint set.

The perturbations will also preserve incentive compatibility under the relaxed problem in the earlier period t. Given the iid assumption, the discounted value of the $\Delta(\delta)$ perturbation at time t is $\beta\delta\pi_\Theta\left(\theta_{t+1} > \theta_{t+1}^n\right)$ to *both* the agent of type θ_t and the agent whose type is one higher and chooses to mimic (that is, δ units of extra utility received if and only if one's type exceeds θ_{t+1}^n). By perturbing the t-dated utility received by both truth-teller and mimicker by $-\beta\delta\pi_\Theta\left(\theta_{t+1} > \theta_{t+1}^n\right)$ units, we ensure the impact on the net present value of utility is zero for both. Hence truth-telling will remain an optimal strategy for both.

The present value (from the perspective of time t) of the cost to the policymaker of applying the $\Delta_{-1}(\delta)$ and $\Delta(\delta)$ perturbations is:

$$
\begin{aligned}
& R_t^{-1}\left\{\pi_\Theta\left(\theta_{t+1}^n\right)\left[\varphi^c\left(\theta_{t+1}^n,\delta\right)-\varphi^y\left(\theta_{t+1}^n,\delta\right)\right]\right. \\
& \left.+\sum_{m=n+1}^{N}\pi_\Theta\left(\theta_{t+1}^m\right)\left[\phi^c\left(\theta_{t+1}^m,\delta\right)-\phi^y\left(\theta_{t+1}^m,\delta\right)\right]\right\} \\
& +\phi^c\left(\theta_t,-\beta\delta\pi_\Theta\left(\theta_{t+1}>\theta_{t+1}^n\right)\right)-\phi^y\left(\theta_t,-\beta\delta\pi_\Theta\left(\theta_{t+1}>\theta_{t+1}^n\right)\right)
\end{aligned}
$$

(where we have suppressed dependence on equilibrium allocations in the ϕ and φ functions to ease the notation). The perturbations have zero net impact on the present value of utility assessed from the perspective of the inital time period, so a necessary condition for optimality is that the derivative of this surplus with respect to δ should equal zero when $\delta = 0$:

$$\begin{aligned} & R_t^{-1} \left\{ \pi_\Theta \left(\theta_{t+1}^n\right) \left[\varphi_2^c \left(\theta_{t+1}^n, 0\right) - \varphi_2^y \left(\theta_{t+1}^n, 0\right)\right] \right. \\ & \left. + \sum_{m=n+1}^{N} \pi_\Theta \left(\theta_{t+1}^m\right) \left[\phi_2^c \left(\theta_{t+1}^m, 0\right) - \phi_2^y \left(\theta_{t+1}^m, 0\right)\right] \right\} \\ & -\beta \pi_\Theta \left(\theta_{t+1} > \theta_{t+1}^n\right) \left[\phi_2^c \left(\theta_t, 0\right) - \phi^y \left(\theta_t, 0\right)\right] \\ = \ & 0 \end{aligned} \tag{A.4}$$

Again, differentiating the equations defining ϕ we can show:

$$\phi_2^c (\theta, 0) - \phi_2^y (\theta, 0) = \frac{1 - \alpha(\theta)}{u_c(\theta) + u_y(\theta)\alpha(\theta)} \tag{A.5}$$

and total differentiation of (A.1) and (A.2) gives:

$$\begin{aligned} \varphi_2^c (\theta, 0) - \varphi_2^y (\theta, 0) &= \frac{1 + \frac{u_c(\theta)}{u_y(\theta)}}{u_c\left(\widehat{\theta}; \theta'\right) - u_y\left(\widehat{\theta}; \theta'\right) \frac{u_c(\theta)}{u_y(\theta)}} \\ &= -\frac{\tau(\theta)}{u_c\left(\widehat{\theta}; \theta'\right)\left(1 - \tau(\theta)\right) + u_y\left(\widehat{\theta}; \theta'\right)} \end{aligned} \tag{A.6}$$

where $\theta' = \min\left\{\theta'' \in \Theta : \theta'' > \theta\right\}$. Using these results in the optimality condition we have:

$$\begin{aligned} & -\pi_\Theta\left(\theta_{t+1}^n\right) \frac{\tau\left(\theta_{t+1}^n\right)}{u_c\left(\widehat{\theta}_{t+1}^n; \theta_{t+1}^{n+1}\right)\left(1 - \tau\left(\theta_{t+1}^n\right)\right) + u_y\left(\widehat{\theta}_{t+1}^n; \theta_{t+1}^{n+1}\right)} \\ & + \sum_{m=n+1}^{N} \pi_\Theta\left(\theta_{t+1}^m\right) \frac{1 - \alpha\left(\theta_{t+1}^m\right)}{u_c\left(\theta_{t+1}^m\right) + u_y\left(\theta_{t+1}^m\right)\alpha\left(\theta_{t+1}^m\right)} \\ = \ & \beta R_{t+1} \pi_\Theta\left(\theta_{t+1} > \theta_{t+1}^n\right) \frac{1 - \alpha(\theta_t)}{u_c(\theta_t) + u_y(\theta_t)\alpha(\theta_t)} \end{aligned} \tag{A.7}$$

QED.

A.4 Proof of Lemma 1

Our focus is restricted to remaining within the 'relaxed' constraint set, so we need only show that it is possible to change the consumption and output levels of each agent in such a way that utilities change in the manner described in the Lemma, and 'downwards' incentive compatibility restrictions remain satisfied for all δ in an open neighbourhood of 0. This requires that the following two conditions are satisfied at $t+1$ for all $n \in \{1, ..., N\}$:

$$u\left(c_{n,t+1}^{*} + \delta_{n}^{c}(\delta), y_{n,t+1}^{*} + \delta_{n}^{y}(\delta); \theta_{t+1}^{n}\right) = u\left(c_{n,t+1}^{*}, y_{n,t+1}^{*}; \theta_{t+1}^{n}\right) + \nu_{n}\delta \quad \text{(A.8)}$$

$$u\left(c_{n,t+1}^{*} + \delta_{n}^{c}(\delta), y_{n,t+1}^{*} + \delta_{n}^{y}(\delta); \theta_{t+1}^{n+1}\right) = u\left(c_{n,t+1}^{*}, y_{n,t+1}^{*}; \theta_{t+1}^{n+1}\right) + \nu_{n+1}\delta \quad \text{(A.9)}$$

where $\delta_{n}^{c}(\delta)$ and $\delta_{n}^{y}(\delta)$ are the perturbations to the nth agent's consumption and output levels respectively. For the Nth agent we just need:

$$u\left(c_{N,t+1}^{*} + \delta_{N}^{c}(\delta), y_{N,t+1}^{*}; \theta_{t+1}^{N}\right) = u\left(c_{N,t+1}^{*}, y_{N,t+1}^{*}; \theta_{t+1}^{N}\right) + \nu_{N}\delta \quad \text{(A.10)}$$

and we normalise $\delta_{N}^{y}(\delta) = 0$.[89]

Equations (A.8) and (A.10) here are just stating that the truth-telling agent should be moved onto a within-period indifference curve consistent with the perturbed utility level obtaining, whilst condition (A.9) states that the specific perturbed allocation should be at a point on this indifference curve such that the change in the utility of a mimicking higher-type agent is equal to the change in that higher-type agent's truth-telling utility. By the single-crossing condition higher-type agents see their utility change monotonically through movements along the indifference curve of a lower-type agent, so for small enough δ these equations must solve for unique values of $\delta_{n}^{c}(\delta)$ and $\delta_{n}^{y}(\delta)$ for all n

[89] This is analogous to the normalisation $\phi^{y}(\theta, k; c^{*}, y^{*}) = 0$ in equation (5.11).

– appealing to the interiority of the solution. (The limit on the magnitude of δ comes from the fact that there is a mimimum level of utility a mimicking agent can obtain along a given lower-type agent's indifference curve.) These values will preserve incentive compatibility at $t+1$. The impact of the $t+1$ perturbations on discounted expected utility from the perspective of time t for an agent who has reported θ^t is to increase it by an amount $\beta\nu'\pi_{\Theta}^{vec}\delta$. If $\nu'\pi_{\Theta}^{vec}=0$ then we are done (confirming the last claim in the Lemma). Otherwise, to preserve 'downward' incentive compatibility at t (and earlier) we must reduce within-period utility in that period by an equal amount, in a manner that has an equal impact on the agent whose true type is θ_t and a mimicker whose type is one higher. In the iid case this can be done through the perturbation

$$\left[\phi^c\left(\theta_t,-\beta\nu'\pi_{\Theta}^{vec}\delta\right),\phi^y\left(\theta_t,-\beta\nu'\pi_{\Theta}^{vec}\delta\right)\right]$$

as already established.

A.5 Proof of Lemma 2

Note first from the definition of γ that $\gamma_n=\nu_n-\nu_{n-1}$ for all $n\in\{1,...,N\}$, where we define $\nu_0=0$. The marginal resource cost at $t+1$ of the perturbation set out in Lemma 1 is:

$$\pi_{\Theta}\left(\theta^t\right)\sum_{n=1}^{N}\pi_{\Theta}\left(\theta_{t+1}^n|\theta_t\right)\left(\frac{d\delta_n^c\left(0\right)}{d\delta}-\frac{d\delta_n^y\left(0\right)}{d\delta}\right) \tag{A.11}$$

where the $\delta_n^c\left(\delta\right)$ and $\delta_n^y\left(\delta\right)$ functions satisfy restrictions (A.8) to (A.10), and $\delta_N^y\left(\delta\right)$ is again set to zero. Totally differentiating those restrictions, one can show:

$$\frac{d\delta_n^c\left(0\right)}{d\delta}=\frac{\frac{u_y\left(\theta_{t+1}^n\right)}{u_c\left(\theta_{t+1}^n\right)}\frac{\nu_{n+1}}{u_c\left(\widehat{\theta}_{t+1}^n;\theta_{t+1}^{n+1}\right)}-\frac{u_y\left(\widehat{\theta}_{t+1}^n;\theta_{t+1}^{n+1}\right)}{u_c\left(\widehat{\theta}_{t+1}^n;\theta_{t+1}^{n+1}\right)}\frac{\nu_n}{u_c\left(\theta_{t+1}^n\right)}}{\frac{u_y\left(\theta_{t+1}^n\right)}{u_c\left(\theta_{t+1}^n\right)}-\frac{u_y\left(\widehat{\theta}_{t+1}^n;\theta_{t+1}^{n+1}\right)}{u_c\left(\widehat{\theta}_{t+1}^n;\theta_{t+1}^{n+1}\right)}} \tag{A.12}$$

$$\frac{d\delta_n^y(0)}{d\delta} = \frac{\frac{\nu_n}{u_c\left(\theta_{t+1}^n\right)} - \frac{\nu_{n+1}}{u_c\left(\widehat{\theta}_{t+1}^n;\theta_{t+1}^{n+1}\right)}}{\frac{u_y\left(\theta_{t+1}^n\right)}{u_c\left(\theta_{t+1}^n\right)} - \frac{u_y\left(\widehat{\theta}_{t+1}^n;\theta_{t+1}^{n+1}\right)}{u_c\left(\widehat{\theta}_{t+1}^n;\theta_{t+1}^{n+1}\right)}} \tag{A.13}$$

for $n \in \{1, ..., N-1\}$, and

$$\frac{d\delta_N^c(0)}{d\delta} = \frac{\nu_N}{u_c\left(\theta_{t+1}^N\right)} \tag{A.14}$$

With some manipulation it is then possible to show for all $n \in \{1, ..., N-1\}$:

$$\begin{aligned} \frac{d\delta_n^c(0)}{d\delta} - \frac{d\delta_n^y(0)}{d\delta} &= \nu_n \frac{1-\alpha\left(\theta_{t+1}^n\right)}{u_c\left(\theta_{t+1}^n\right) + \alpha\left(\theta_{t+1}^n\right) u_y\left(\theta_{t+1}^n\right)} \\ &\quad - \left(\nu_{n+1} - \nu_n\right) \frac{\tau\left(\theta_{t+1}^n\right)}{u_c\left(\widehat{\theta}_{t+1}^n;\theta_{t+1}^{n+1}\right)\left(1-\tau\left(\theta_{t+1}^n\right)\right) + u_y\left(\widehat{\theta}_{t+1}^n;\theta_{t+1}^{n+1}\right)} \\ &= \nu_n \left(\phi_2^c\left(\theta_{t+1}^n, 0\right) - \phi_2^y\left(\theta_{t+1}^n, 0\right)\right) \\ &\quad + \left(\nu_{n+1} - \nu_n\right)\left(\varphi_2^c\left(\theta_{t+1}^n, 0\right) - \varphi_2^y\left(\theta_{t+1}^n, 0\right)\right) \end{aligned} \tag{A.15}$$

and

$$\begin{aligned} \frac{d\delta_N^c(0)}{d\delta} - \frac{d\delta_N^y(0)}{d\delta} &= \frac{d\delta_N^c(0)}{d\delta} \\ &= \frac{\nu_N}{u_c\left(\theta_{t+1}^N\right)} \\ &= \nu_N \left(\phi_2^c\left(\theta_{t+1}^N, 0\right) - \phi_2^y\left(\theta_{t+1}^N, 0\right)\right) \end{aligned} \tag{A.16}$$

where we apply the earlier definitions of the ϕ^c, ϕ^y, φ^c and φ^y functions. Hence:

$$\begin{aligned}
&\sum_{n=1}^{N} \pi_\Theta \left(\theta_{t+1}^n | \theta_t\right) \left(\frac{d\delta_n^c(0)}{d\delta} - \frac{d\delta_n^y(0)}{d\delta} \right) \\
= &\sum_{n=1}^{N} \pi_\Theta \left(\theta_{t+1}^n | \theta_t\right) \left[\nu_n \left(\phi_2^c \left(\theta_{t+1}^n, 0\right) - \phi_2^y \left(\theta_{t+1}^n, 0\right) \right) \right. \\
&\left. + \left(\nu_{n+1} - \nu_n\right) \left(\varphi_2^c \left(\theta_{t+1}^n, 0\right) - \varphi_2^y \left(\theta_{t+1}^n, 0\right) \right) \right] \\
= &\sum_{n=1}^{N-1} \pi_\Theta \left(\theta_{t+1}^n | \theta_t\right) \left[\sum_{m=1}^{n} \gamma_m \left(\phi_2^c \left(\theta_{t+1}^n, 0\right) - \phi_2^y \left(\theta_{t+1}^n, 0\right) \right) \right. \\
&\left. + \gamma_{n+1} \left(\varphi_2^c \left(\theta_{t+1}^n, 0\right) - \varphi_2^y \left(\theta_{t+1}^n, 0\right) \right) \right] \\
&+ \pi_\Theta \left(\theta_{t+1}^N | \theta_t\right) \sum_{m=1}^{N} \gamma_m \left(\phi_2^c \left(\theta_{t+1}^N, 0\right) - \phi_2^y \left(\theta_{t+1}^N, 0\right) \right)
\end{aligned} \tag{A.17}$$

This last expression can equivalently be written in matrix form:

$$\left(\pi_\Theta^{vec}\right)' \left[\sum_{n=1}^{N} \gamma_n \Delta^{n\prime}(0) \right] k \tag{A.18}$$

The first part of the result follows. For the second we need to show additionally that it is not possible to move in any dimension *not* described by a $\Delta^{n\prime}(0)$ matrix and preserve incentive compatibility for the relaxed problem. But the only degree of freedom we have to vary the above changes is in relaxing the normalisation that $\delta_N^y(\delta) = 0$, whilst still satisfying

$$u\left(c_{N,t+1}^* + \delta_N^c(\delta), y_{N,t+1}^* + \delta_N^y(\delta); \theta_{t+1}^N\right) = u\left(c_{N,t+1}^*, y_{N,t+1}^*; \theta_{t+1}^N\right) + \nu_N \delta \tag{A.19}$$

This gives:

$$u_c\left(\theta_{t+1}^N\right) \frac{d\delta_N^c(0)}{d\delta} + u_y\left(\theta_{t+1}^N\right) \frac{d\delta_N^y(0)}{d\delta} = \nu_N \tag{A.20}$$

But at the top $u_c\left(\theta_{t+1}^N\right) = -u_y\left(\theta_{t+1}^N\right)$, so:

$$\begin{aligned}
\frac{d\delta_N^c(0)}{d\delta} - \frac{d\delta_N^y(0)}{d\delta} &= \frac{\nu_N}{u_c\left(\theta_{t+1}^N\right)} \\
&= \nu_N\left(\phi_2^c\left(\theta_{t+1}^N, 0\right) - \phi_2^y\left(\theta_{t+1}^N, 0\right)\right) \\
&= \nu_N\left(\phi_2^c\left(\theta_{t+1}^N, 0\right) - \phi_2^y\left(\theta_{t+1}^N, 0\right)\right) \\
&\quad +\kappa\left(\varphi_2^c\left(\theta_{t+1}^N, 0\right) - \varphi_2^y\left(\theta_{t+1}^N, 0\right)\right)
\end{aligned} \tag{A.21}$$

for arbitrary κ. So if $\frac{\delta_N^y(0)}{d\delta}$ differs from zero then $\frac{\delta_N^c(0)}{d\delta}$ can differ from $\frac{\nu_N}{u_c\left(\theta_{t+1}^N\right)}$, but only in a manner that raises no net resources – which is equivalent to a movement along the 'top' indifference curve, given that the optimum involves no distortion at the top. So the only additional dimension in which outcomes can be perturbed at the margin is that described by $\Delta^{N+1\prime}(0)$. Hence the complete set of $\Delta^{n\prime}(0)$ matrices does indeed span the relevant space.

A.6 Proof of Lemma 4

Again, the Lemma requires us to focus only on the need to ensure 'downwards' incentive compatibility continues to hold locally at t and $t+1$. The latter is simpler: it requires that the following conditions are satisfied for agents with the relevant reporting history for all $m \in \{1, ..., N\}$:

$$u\left(c_{m,t+1}^* + \delta_{m,t+1}^c(\delta), y_{m,t+1}^* + \delta_{m,t+1}^y(\delta); \theta_{t+1}^m\right) = u\left(c_{m,t+1}^*, y_{m,t+1}^*; \theta_{t+1}^m\right) + \nu_m\delta \tag{A.22}$$

$$u\left(c_{m,t+1}^* + \delta_{m,t+1}^c(\delta), y_{m,t+1}^* + \delta_{m,t+1}^y(\delta); \theta_{t+1}^{m+1}\right) = u\left(c_{m,t+1}^*, y_{m,t+1}^*; \theta_{t+1}^{m+1}\right) + \nu_{m+1}\delta \tag{A.23}$$

where $\delta_{m,t+1}^c(\delta)$ and $\delta_{m,t+1}^y(\delta)$ are the perturbations to the mth agent's consumption and output levels respectively. For the Nth agent we just need:

$$u\left(c_{N,t+1}^* + \delta_{N,t+1}^c(\delta), y_{N,t+1}^*; \theta_{t+1}^N\right) = u\left(c_{N,t+1}^*, y_{N,t+1}^*; \theta_{t+1}^N\right) + \nu_N\delta \tag{A.24}$$

and we normalise $\delta^{y}_{N,t+1}(\delta) = 0$.

The proof of Lemma 1 shows that these conditions can indeed be satisfied by appropriate choice of $\delta^{c}_{m,t+1}(\delta)$ and $\delta^{y}_{m,t+1}(\delta)$ schedules, given an interior optimum. There remains the problem of incentive compatibility (under the relaxed problem) at t. From the perspective of that time period the $t+1$ perturbations are increasing expected utility for potential mimickers by $\beta\delta$ units, whilst leaving that of truth-tellers constant. To offset this effect we need to move along the indifference curve of the nth agent at t to such an extent that a mimicker's utility is reduced by an offsetting amount (whilst, by definition, leaving the utility of a truth-teller unaffected in this period also). That requires $\delta^{c}_{n,t}(\delta)$ and $\delta^{y}_{n,t}(\delta)$ schedules that satisfy:

$$u\left(c^{*}_{n,t} + \delta^{c}_{n,t}(\delta), y^{*}_{n,t} + \delta^{y}_{n,t}(\delta); \theta^{n}_{t}\right) = u\left(c^{*}_{n,t}, y^{*}_{n,t}; \theta^{n}_{t}\right) \tag{A.25}$$

$$u\left(c^{*}_{n,t} + \delta^{c}_{n,t}(\delta), y^{*}_{n,t} + \delta^{y}_{n,t}(\delta); \theta^{n+1}_{t}\right) = u\left(c^{*}_{n,t}, y^{*}_{n,t}; \theta^{n+1}_{t}\right) - \beta\delta \tag{A.26}$$

Again, by the single crossing condition the utility of the agent of type θ^{n+1}_{t} changes monotonically as one moves along a lower-type agent's indifference curve, so for small enough δ in an open neighbourhood of $\delta = 0$ this is always possible – with a limit provided by the fact that there is a mimimum to the utility that mimickers can obtain on the given lower-type indifference curve.

A.7 Proof of Lemma 6

It remains to establish the result for the case in which consumption and labour supply are Edgeworth complements (in which case $\alpha(\theta_t) > 0$) and productivities follow an iid process. In order to put a zero lower bound on the marginal cost of utility provision in this case we need to verify that $\alpha(\theta_t) < 1$ – that is, that the marginal cost of incentive-compatible utility provision never turns negative under an optimal plan. Suppose instead that $\alpha(\theta_t) \geq 1$ were to hold for some θ_t and a given report history. We argue that in this situation it is always possible for the policymaker to generate surplus resources at the margin, whilst preserving incentive compatibility – contradicting optimality.

If $\alpha(\theta_t) \geq 1$ then the generalised inverse Euler equation implies we must also have:

$$\sum_{\theta_{t+1}\in\Theta} \pi_{\Theta}(\theta_{t+1}|\theta_t)\left[\frac{1-\alpha(\theta_{t+1})}{u_c(\theta_{t+1})+u_y(\theta_{t+1})\alpha(\theta_{t+1})}\right] \geq 1 \tag{A.27}$$

With iid shocks it is always possible for us to find an $(N+1)\times 1$ vector γ that satisfies the following:

$$\left[\sum_{n=1}^{N+1} \gamma_n \Delta^{n\prime}(0)\right] = \widetilde{\Delta} \tag{A.28}$$

for some $N \times 2$ $\widetilde{\Delta}$ matrix whose first column is an $N \times 1$ vector of strictly positive scalars and whose second is an $N \times 1$ vector of zeros. To see this, note that the marginal vector $\Delta^{n\prime}(0)$ has zeros in all rows up to the $(n-2)$th, and an entry in the $(n-1)$th row that is linearly independent of the $(n-1)$th row in all of the other $\Delta^{m\prime}(0)$ matrices. So we can choose the entries of γ by first finding an additive combination of $\Delta^{1\prime}(0)$ and $\Delta^{2\prime}(0)$ sufficient to increase the consumption of the agent of type θ^1_{t+1} without changing that agent's output requirements, then add to this sufficient units of $\Delta^{3\prime}(0)$ for the output level of the agent of type θ^2_{t+1} to remain unchanged (this perturbation has no impact on type θ^1_{t+1}), then sufficient units of $\Delta^{4\prime}(0)$ for the output level of the agent of type θ^3_{t+1} to remain unchanged, and so on. Since consumption is increasing for the agent of type θ^1_{t+1}, it must also increase for all higher-type agents – otherwise we could not continue to satisfy incentive compatibility. This $\widetilde{\Delta}$ perturbation increases the *ex ante* expected utility level of an agent with the relevant reporting history, so it will be possible to leave that expected utility level unchanged (preserving incentive compatibility at t) by a linear combination of $\Delta^{1\prime}(0)$ and $\widetilde{\Delta}$, with a positive coefficient on the former and a negative on the latter. But if $\widetilde{\Delta}$ is providing utility through consumption increments alone it must come at a positive cost, whilst we have from (A.27) that the cost of $\Delta^{1\prime}(0)$ is negative. Hence we raise a surplus, contradicting optimality.[90]

[90] Notice that this argument cannot be applied in the case of non-iid productivity processes, since the perturbation

A.8 Proof of Proposition 12

We know Doob's convergence theorem applies to the non-negative martingale $\frac{1-\alpha(\theta_t)}{u_c(\theta_t)+u_y(\theta_t)\alpha(\theta_t)}$, so need only show that it is not possible for this object to converge to any non-zero value. The following Lemma is useful:

Lemma 7 $\frac{\tau(\theta_t^n)}{u_c\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)(1-\tau(\theta_t^n))+u_y\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)} \overset{a.s.}{\to} 0$ *holds under an optimal plan that solves the restricted problem.*

Proof. In the iid case this follows directly from equation (5.28):

$$\begin{aligned} & \lim_{t\to\infty}\left[-\pi_\Theta\left(\theta_{t+1}^n|\theta_t\right)\frac{\tau\left(\theta_{t+1}^n\right)}{u_c\left(\widehat{\theta}_{t+1}^n;\theta_{t+1}^{n+1}\right)\left(1-\tau\left(\theta_{t+1}^n\right)\right)+u_y\left(\widehat{\theta}_{t+1}^n;\theta_{t+1}^{n+1}\right)}\right] \qquad \text{(A.29)} \\ = & -\sum_{m=n+1}^{N}\pi_\Theta\left(\theta_{t+1}^m|\theta_t\right)\lim_{t\to\infty}\left[\frac{1-\alpha\left(\theta_{t+1}^m\right)}{u_c\left(\theta_{t+1}^m\right)+u_y\left(\theta_{t+1}^m\right)\alpha\left(\theta_{t+1}^m\right)}\right] \\ & +\pi_\Theta\left(\theta_{t+1}>\theta_{t+1}^n|\theta_t\right)\sum_{\theta_{t+1}\in\Theta}\pi_\Theta\left(\theta_{t+1}|\theta_t\right)\lim_{t\to\infty}\left[\frac{1-\alpha\left(\theta_{t+1}\right)}{u_c\left(\theta_{t+1}\right)+u_y\left(\theta_{t+1}\right)\alpha\left(\theta_{t+1}\right)}\right] \\ = & \ 0 \end{aligned}$$

In the Markov case we know that equation (5.28) must hold in periods immediately following those in which $\theta=\theta^N$, and so if one indexes by T the (infinite) set of periods in which this is the case, and denotes by $t(T)$ the (conventional) time period corresponding to the Tth occasion on which $\theta=\theta^N$

operating on consumption levels alone does not generate a *uniform* level of incremental utility provision across $t+1$ types, and thus will have differential effects on the expected utility levels of mimickers and truth-tellers at time t. One could generalise to consider the complete set of perturbations guaranteed to increase utility for all agents at $t+1$ at a positive cost – that is, movements that increase consumption by a greater amount than output at the margin for all agents, whilst simultaneously increasing utility. But even this set is not sufficiently large to prove the result: the movements it permts do not span an entire half-space in the (c_{t+1},y_{t+1}) plane for each agent.

has obtained along the given sample path, we must have:

$$\begin{aligned}
& \lim_{T\to\infty}\Bigg[-\pi_\Theta\left(\theta^n_{t(T)+1}|\theta_{t(T)}\right) \\
& \cdot\frac{\tau\left(\theta^n_{t(T)+1}\right)}{u_c\left(\widehat{\theta}^n_{t(T)+1};\theta^{n+1}_{t(T)+1}\right)\left(1-\tau\left(\theta^n_{t(T)+1}\right)\right)+u_y\left(\widehat{\theta}^n_{t(T)+1};\theta^{n+1}_{t(T)+1}\right)}\Bigg] \\
= & -\sum_{m=n+1}^{N}\pi_\Theta\left(\theta^m_{t(T)+1}|\theta_{t(T)}\right)\lim_{T\to\infty}\left[\frac{1-\alpha\left(\theta^m_{t(T)+1}\right)}{u_c\left(\theta^m_{t(T)+1}\right)+u_y\left(\theta^m_{t(T)+1}\right)\alpha\left(\theta^m_{t(T)+1}\right)}\right] \\
& +\pi_\Theta\left(\theta_{t(T)+1}>\theta^n_{t(T)+1}|\theta_{t(T)}\right) \\
& \cdot\sum_{\theta_{t(T)+1}\in\Theta}\pi_\Theta\left(\theta_{t(T)+1}|\theta_{t(T)}\right)\lim_{T\to\infty}\left[\frac{1-\alpha\left(\theta_{t(T)+1}\right)}{u_c\left(\theta_{t(T)+1}\right)+u_y\left(\theta_{t(T)+1}\right)\alpha\left(\theta_{t(T)+1}\right)}\right] \\
= & \ 0
\end{aligned} \tag{A.30}$$

But if

$$\frac{\tau\left(\theta^n_{t(T)+1}\right)}{u_c\left(\widehat{\theta}^n_{t(T)+1};\theta^{n+1}_{t(T)+1}\right)\left(1-\tau\left(\theta^n_{t(T)+1}\right)\right)+u_y\left(\widehat{\theta}^n_{t(T)+1};\theta^{n+1}_{t(T)+1}\right)}=0$$

holds at the limit as T becomes large then we must also, at the same limit, have an identical set of zero restrictions in period $t(T)+2$, by equations (6.5) and (6.4). By induction this can then be extended to period $t(T)+n$ for all $n>1$, and the result follows. ■

This Lemma implies two alternatives: either

$$\tau\left(\theta^n_t\right)\overset{a.s.}{\to}0$$

or

$$u_c\left(\widehat{\theta}^n_t;\theta^{n+1}_t\right)\left(1-\tau\left(\theta^n_t\right)\right)+u_y\left(\widehat{\theta}^n_t;\theta^{n+1}_t\right)\overset{a.s.}{\to}\infty$$

Suppose the latter were true. By equation (7.4) we have:

$$\begin{aligned} u_c\left(\theta_t^n\right) + u_y\left(\theta_t^n\right)\alpha\left(\theta_t^n\right) &= \frac{u_c\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right) - u_y\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)\frac{u_c\left(\theta_t^n\right)}{u_y\left(\theta_t^n\right)}}{1-\frac{u_y\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)}{u_y\left(\theta_t^n\right)}} \\ &> u_c\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right) - u_y\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)\frac{u_c\left(\theta_t^n\right)}{u_y\left(\theta_t^n\right)} \\ &= u_c\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right) - u_y\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)\frac{1}{\left(1-\tau\left(\theta_t^n\right)\right)} \end{aligned} \tag{A.31}$$

If

$$u_c\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)\left(1-\tau\left(\theta_t^n\right)\right) + u_y\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right) \overset{a.s.}{\to} \infty$$

then

$$u_c\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right) - u_y\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)\frac{1}{\left(1-\tau\left(\theta_t^n\right)\right)} \overset{a.s.}{\to} \infty$$

must also hold, since $\left(1-\tau\left(\theta_t^n\right)\right) \in [0,1]$ follows from the definition of τ and Proposition 8. Hence we must also have

$$u_c\left(\theta_t^n\right) + u_y\left(\theta_t^n\right)\alpha\left(\theta_t^n\right) \overset{a.s.}{\to} \infty$$

This in turn implies $\frac{1-\alpha\left(\theta_t^n\right)}{u_c\left(\theta_t^n\right)+u_y\left(\theta_t^n\right)\alpha\left(\theta_t^n\right)}$ can only converge to a non-zero limit if $\left|\alpha\left(\theta_t\right)\right|$ is itself always infinite at that limit. But since we know $\alpha\left(\theta_t\right) = 0$ when $\theta_t = \theta^N$ we can rule that out.

The alternative is that $\tau\left(\theta_t^n\right) \overset{a.s.}{\to} 0$. In this case we have $u_c\left(\theta_t^n\right) = -u_y\left(\theta_t^n\right)$ at the limit, and so

$$\frac{1-\alpha\left(\theta_t^n\right)}{u_c\left(\theta_t^n\right) + u_y\left(\theta_t^n\right)\alpha\left(\theta_t^n\right)} = \frac{1}{u_c\left(\theta_t^n\right)}$$

Hence the inverse of the marginal utility of consumption must be converging to a common value for all agents. But since $u_c\left(\theta_t^n\right) = -u_y\left(\theta_t^n\right)$ the marginal disutility of production must also be converging to the *same* value across agents. Suppose this were a finite value. We have shown when analysing the first-best that if u_c is common across types and $u_c = -u_y$ holds then utility must be *decreasing* in type.

This is clearly inconsistent with incentive compatibility, which is enough to rule out $\frac{1-\alpha(\theta_t^n)}{u_c(\theta_t^n)+u_y(\theta_t^n)\alpha(\theta_t^n)}$ converging to a non-zero value in this case too. This completes the proof.

B Appendix to Part III

B.1 Derivation of equation 5.7

The forward-looking policymaker's objective for this problem can be written as:

$$\max_{c(h_\theta,B),y(h_\theta,B)} (1-\beta)^{-1} \int_{\Theta^\infty\times\mathbb{R}} \sum_{\theta\in\Theta} \{u(h_\theta,B;\theta)\,\pi_\Theta(\theta|h_\theta) \tag{B.1}$$

$$+ \sum_{\theta\in\Theta\setminus\underline{\theta}} \lambda_{IC}(h_\theta,B;\theta)\left[V(h_\theta,B;\theta) - \widehat{V}(h_\theta,B;\theta)\right]$$

$$+ \sum_{\theta\in\Theta} \lambda_V^0(h_\theta,B;\theta)\,\pi_\Theta(\theta|h_\theta)\left[V(h_\theta,B;\theta) - u(h_\theta,B;\theta)\right]$$

$$+ \sum_{\theta\in\Theta\setminus\underline{\theta}} \lambda_{\widehat{V}}^0(h_\theta,B;\theta)\,\pi_\Theta\left(\theta'|h_\theta\right)\left[\widehat{V}(h_\theta,B;\theta) - \widehat{u}(h_\theta,B;\theta)\right]$$

$$- \sum_{\theta\in\Theta} \lambda_V^1(h_\theta,B;\theta)\,\beta\pi_\Theta(\theta|h_\theta)\,V(h_\theta,B;\theta)$$

$$- \sum_{\theta\in\Theta} \lambda_{\widehat{V}}^1(h_\theta,B;\theta)\,\beta\pi_\Theta\left(\theta|h'_\theta\right)V(h_\theta,B;\theta)$$

$$+\lambda_R(B)\left[RB + \sum_{\theta\in\Theta}\left[y(h_\theta,B;\theta) - c(h_\theta,B;\theta)\right]\pi_\Theta(\theta|h_\theta) - B(B)\right]\Bigg\}\,dF^0(h_\theta,B)$$

where again we keep the consumption and output arguments of u and $\widehat{u}$ implicit, but reintroduce dependence on the aggregate asset stock to retain consistency with the notation of the general problem.[91] The superscripts on the λ_V and $\lambda_{\widehat{V}}$ multipliers are equivalent to the index s in the general objective (2.20), which in this case only runs from 0 to 1. The history h'_θ is defined as that obtained by replacing the terminal entry in h_θ with the element one higher in Θ, and $\theta' = \max\left\{\widetilde{\theta}\in\Theta : \widetilde{\theta} < \theta\right\}$ for the given θ. If the terminal entry in h_θ is $\overline{\theta}$ then we let $h'_\theta = h_\theta$, and normalise $\lambda_{\widehat{V}}^1(h_\theta,B;\theta) = 0$.

To determine optimal forward-looking policy we will consider the marginal effects of joint perturbations to consumption and output together, akin to those constructed (and explored in more detail)

[91] $\lambda_R(B)$ could always be denoted $\lambda_R(h_\theta,B)$, but this shadow value of resources will be invariant in shock histories given the absence of aggregate uncertainty.

in Part II. Since u and $\widehat{u}$ are in turn determined entirely by agents' consumption and output allocations (for given types), it will be necessary to consider simultaneously the marginal effects of joint consumption-output perturbations on (B.1) through changes to all *four* of the objects c, y, u, and $\widehat{u}$. As in Part II, we want to make sure that there are equal effects on both truthful reporters and mimickers of any given type when outcomes are perturbed. This means that for every unit by which we wish to increase the within-period utility of an agent of history h_θ and current type θ at the margin we raise consumption by an amount:

$$\frac{dc\left(h_\theta, B; \theta\right)}{du\left(h_\theta, B; \theta\right)} = \frac{1}{u_c\left(h_\theta, B; \theta\right) + u_y\left(h_\theta, B; \theta\right)\alpha\left(h_\theta, B; \theta\right)} \tag{B.2}$$

and raise output by an amount:

$$\frac{dy\left(h_\theta, B; \theta\right)}{du\left(h_\theta, B; \theta\right)} = \frac{\alpha\left(h_\theta, B; \theta\right)}{u_c\left(h_\theta, B; \theta\right) + u_y\left(h_\theta, B; \theta\right)\alpha\left(h_\theta, B; \theta\right)} \tag{B.3}$$

where $\alpha\left(h_\theta; \theta\right)$ is again defined as the increase in output required to ensure a would-be mimicker's within-period utility is increased by the perturbation by an identical amount as the agent whose type truly corresponds to θ:

$$\alpha\left(h_\theta, B; \theta\right) = \frac{u_c\left(h_\theta, B; \theta\right) - \widehat{u}_c\left(h_\theta, B; \theta'\right)}{\widehat{u}_y\left(h_\theta, B; \theta'\right) - u_y\left(h_\theta, B; \theta\right)} \tag{B.4}$$

where $\theta' = \min\left\{\widetilde{\theta} \in \Theta : \widetilde{\theta} > \theta\right\}$, and we denote marginal utilities by subscripts as usual (If $\theta = \overline{\theta}$ we simply define $\alpha\left(h_\theta, B; \theta\right) = 0$.)

The first-order condition for the forward-looking policymaker associated with a perturbation that increases $u\left(h_\theta, B; \theta\right)$ by one unit at the margin through consumption-output changes in these proportions is:

$$1 - \lambda_R\left(B\right) MC\left(h_\theta, B; \theta\right) - \lambda_V^0\left(h_\theta, B; \theta\right) - \lambda_{\widehat{V}}^0\left(h_\theta, B; \theta'\right) = 0 \tag{B.5}$$

for $\theta \neq \overline{\theta}$ and $\theta' = \min\left\{\widetilde{\theta} \in \Theta : \widetilde{\theta} > \theta\right\}$, where we define the marginal cost of incentive-compatible utility provision as above:

$$MC(h_\theta, B; \theta) \equiv \frac{1 - \alpha(h_\theta, B; \theta)}{u_c(h_\theta, B; \theta) + u_y(h_\theta, B; \theta)\,\alpha(h_\theta, B; \theta)} \tag{B.6}$$

and:

$$1 - \lambda_R(B)\,MC(h_\theta, B; \theta) - \lambda_V^0(h_\theta, B; \theta) = 0 \tag{B.7}$$

for $\theta = \overline{\theta}$.

By considering marginal effects of increasing $V(h_\theta, B; \theta)$ and $\widehat{V}(h_\theta, B; \theta)$ we additionally have, for $\theta \neq \underline{\theta}$:

$$\begin{aligned} & \lambda_{IC}(h_\theta, B; \theta) + \pi_\Theta(\theta|h_\theta)\,\lambda_V^0(h_\theta, B; \theta) \\ & -\beta\pi_\Theta(\theta|h_\theta)\,\lambda_V^1(h_\theta, B; \theta) - \beta\pi_\Theta\left(\theta|h'_\theta\right)\lambda_{\widehat{V}}^1(h_\theta, B; \theta) \\ = \; & 0 \end{aligned} \tag{B.8}$$

for h'_θ as defined above, and

$$-\lambda_{IC}(h_\theta, B; \theta) + \pi_\Theta\left(\theta'|h_\theta\right)\lambda_{\widehat{V}}^0(h_\theta, B; \theta) = 0 \tag{B.9}$$

whilst for $\theta = \underline{\theta}$ there is no incentive compatibility constraint, and we simply have:

$$\begin{aligned} & \pi_\Theta(\theta|h_\theta)\,\lambda_V^0(h_\theta, B; \theta) - \beta\pi_\Theta(\theta|h_\theta)\,\lambda_V^1(h_\theta, B; \theta) \\ & -\beta\pi_\Theta\left(\theta|h'_\theta\right)\lambda_{\widehat{V}}^1(h_\theta, B; \theta) \\ = \; & 0 \end{aligned} \tag{B.10}$$

By taking weighted sums of the preceeding conditions over Θ, we have:

$$\begin{aligned}
&\sum_{\theta \in \Theta} \pi_\Theta(\theta|h_\theta) \lambda_V^0(h_\theta, B; \theta) + \sum_{\theta \in \Theta \backslash \underline{\theta}} \pi_\Theta(\theta'|h_\theta) \lambda_{\widehat{V}}^0(h_\theta, B; \theta) \\
= \; & \beta \sum_{\theta \in \Theta} \pi_\Theta(\theta|h_\theta) \lambda_V^1(h_\theta, B; \theta) + \beta \sum_{\theta \in \Theta} \pi_\Theta(\theta|h'_\theta) \lambda_{\widehat{V}}^1(h_\theta, B; \theta)
\end{aligned} \tag{B.11}$$

for $\theta' = \max\left\{\widetilde{\theta} \in \Theta : \widetilde{\theta} < \theta\right\}$ for the given θ, and where the last term in this expression equals zero if the terminal entry of h_θ is $\overline{\theta}$.

From the definition of a veil of ignorance optimum, part 2 (c), we know:[92]

$$\lambda_V^1\left((h_\theta, \theta), B(B); \theta''\right) = \lambda_V^0(h_\theta, B; \theta) \tag{B.12}$$

must hold whenever the backward-looking policymaker's contingent choice of assets $B(B)$ is a one-to-one mapping, and similarly:

$$\lambda_{\widehat{V}}^1\left((h_\theta, \theta), B(B); \theta''\right) = \lambda_{\widehat{V}}^0\left(h_\theta, B; \theta'\right) \tag{B.13}$$

for $\theta' = \min\left\{\widetilde{\theta} \in \Theta : \widetilde{\theta} > \theta\right\}$. The important point is that in both cases the multipliers on the left-hand side are independent of θ''.

By taking a weighted sum of (B.5) together with (B.7), at input values $((h_\theta, \theta), B(B)) \in \Theta^\infty \times \mathbb{R}$, we can obtain:

$$\begin{aligned}
&\sum_{\theta'' \in \Theta} \pi_\Theta\left(\theta''|h_\theta\right) \lambda_V^0\left((h_\theta, \theta), B(B); \theta''\right) \\
&+ \sum_{\theta'' \in \Theta \backslash \underline{\theta}} \pi_\Theta\left(\theta'|h_\theta\right) \lambda_{\widehat{V}}^0\left((h_\theta, \theta), B(B); \theta''\right) \\
= \; & 1 - \lambda_R(B(B)) \sum_{\theta'' \in \Theta} \pi_\Theta\left(\theta''|h_\theta\right) MC\left((h_\theta, \theta), B(B); \theta''\right)
\end{aligned} \tag{B.14}$$

[92] C.f. equation (2.37).

for $\theta' = \max\left\{\widetilde{\theta} \in \Theta : \widetilde{\theta} < \theta''\right\}$ for the given θ''. Using the three prior expressions in this gives:

$$\begin{aligned} & \beta\left[\lambda_V^0\left(h_\theta, B; \theta\right) + \lambda_{\hat{V}}^0\left(h_\theta, B; \theta'\right)\right] \\ = \; & 1 - \lambda_R\left(B\left(B\right)\right) \sum_{\theta'' \in \Theta} \pi_\Theta\left(\theta'' | h_\theta\right) MC\left(\left(h_\theta, \theta\right), B\left(B\right); \theta''\right) \end{aligned} \tag{B.15}$$

for $\theta \neq \overline{\theta}$, and:

$$\begin{aligned} & \beta\lambda_V^0\left(h_\theta, B; \theta\right) \\ = \; & 1 - \lambda_R\left(B\left(B\right)\right) \sum_{\theta'' \in \Theta} \pi_\Theta\left(\theta'' | h_\theta\right) MC\left(\left(h_\theta, \theta\right), B\left(B\right); \theta''\right) \end{aligned} \tag{B.16}$$

for $\theta = \overline{\theta}$.

Using the previous two results in (B.5) and (B.7) respectively, we have at a veil of ignorance optimum:

$$\begin{aligned} & \beta\left[1 - \lambda_R\left(B\right) MC\left(h_\theta, B; \theta\right)\right] \\ = \; & 1 - \lambda_R\left(B\left(B\right)\right) \sum_{\theta'' \in \Theta} \pi_\Theta\left(\theta'' | h_\theta\right) MC\left(\left(h_\theta, \theta\right), B\left(B\right); \theta''\right) \end{aligned} \tag{B.17}$$

And since $\lambda_R\left(B\right) = \lambda_R\left(B\left(B\right)\right)$ by (5.6) (under the assumption that $R = \beta^{-1}$), we have:

$$\begin{aligned} & \sum_{\theta'' \in \Theta} \pi_\Theta\left(\theta'' | h_\theta\right) \left[MC\left(\left(h_\theta, \theta\right), B\left(B\right); \theta''\right) - \overline{MC}\right] \\ = \; & \beta\left[MC\left(h_\theta, B; \theta\right) - \overline{MC}\right] \end{aligned} \tag{B.18}$$

where $\overline{MC}$ is a constant, equal to $\frac{1}{\lambda_R(B)}$ for the constant stock of assets B that the backward-looking policymaker chooses to hold. Integrating both sides of this equation over all type histories immediately implies that $\overline{MC}$ is the population average of the marginal cost of utility provision in all time periods, accounting for the notation.

B.2 Derivation of equation 5.13

For state (h_θ, B), consider a composite perturbation to the consumption and output of an agent of current type θ, corresponding to a movement along the within-period indifference curve of that agent by an amount just sufficient to increase $\widehat{u}\left(h_\theta, B; \theta'\right)$ by one unit at the margin (where $\theta' = \min\left\{\widetilde{\theta} \in \Theta : \widetilde{\theta} > \theta\right\}$). Using the arguments developed in Part II of this thesis, a necessary condition for the forward-looking policymaker's behaviour to be optimal follows from this:

$$\lambda^0_{\widehat{V}}\left(h_\theta, B; \theta'\right) - \lambda_R\left(B\right) TC\left(h_\theta, B; \theta\right) = 0 \tag{B.19}$$

Using this condition in (B.5), we additionally have:

$$\lambda^0_V\left(h_\theta, B; \theta\right) = 1 - \lambda_R\left(B\right)\left[MC\left(h_\theta, B; \theta\right) + TC\left(h_\theta, B; \theta\right)\right] \tag{B.20}$$

for $\theta \neq \overline{\theta}$.

By taking weighted sums of conditions (B.8) to (B.10) in proportions given by the entries in ν we have the following:[93]

$$\begin{aligned} & \sum_{\theta\in\Theta} \pi_\Theta\left(\theta|h_\theta\right)\nu\left(\theta\right)\lambda^0_V\left(h_\theta, B; \theta\right) + \sum_{\theta\in\Theta\backslash\underline{\theta}} \pi_\Theta\left(\theta'|h_\theta\right)\nu\left(\theta\right)\lambda^0_{\widehat{V}}\left(h_\theta, B; \theta\right) \\ = \;& \beta\sum_{\theta\in\Theta} \pi_\Theta\left(\theta|h_\theta\right)\nu\left(\theta\right)\lambda^1_V\left(h_\theta, B; \theta\right) + \beta\sum_{\theta\in\Theta} \pi_\Theta\left(\theta|h'_\theta\right)\nu\left(\theta\right)\lambda^1_{\widehat{V}}\left(h_\theta, B; \theta\right) \end{aligned} \tag{B.21}$$

for $\theta' = \max\left\{\widetilde{\theta} \in \Theta : \widetilde{\theta} < \theta\right\}$ for the given θ. But by (B.12) and (B.13) we know that $\lambda^1_V\left(h_\theta, B; \theta\right)$

[93] Recall that $\lambda^1_{\widehat{V}}\left(h_\theta, B; \theta\right) = 0$ if the terminal entry in h_θ is the maximal element in Θ.

and $\lambda^1_{\widehat{V}}(h_\theta, B; \theta)$ cannot vary in θ, and thus we must have:

$$\beta \sum_{\theta\in\Theta} \pi_\Theta(\theta|h_\theta)\,\nu(\theta)\,\lambda^1_V(h_\theta, B; \theta) = \beta \sum_{\theta\in\Theta} \pi_\Theta\left(\theta|h'_\theta\right)\nu(\theta)\,\lambda^1_{\widehat{V}}(h_\theta, B; \theta) = 0 \tag{B.22}$$

where we make use of the zero restrictions characterising the ν vector here. Using (B.19) and (B.20) gives:

$$\begin{aligned} & \sum_{\theta\in\Theta} \pi_\Theta(\theta|h_\theta)\,\nu(\theta)\,\lambda^0_V(h_\theta, B; \theta) + \sum_{\theta\in\Theta\backslash\underline{\theta}} \pi_\Theta\left(\theta'|h_\theta\right)\nu(\theta)\,\lambda^0_{\widehat{V}}(h_\theta, B; \theta) \\ = & -\sum_{\theta\in\Theta} \pi_\Theta(\theta|h_\theta)\,\nu(\theta)\,\lambda_R(B)\left[MC(h_\theta, B; \theta) + TC(h_\theta, B; \theta)\right] \\ & + \sum_{\theta\in\Theta\backslash\underline{\theta}} \pi_\Theta\left(\theta'|h_\theta\right)\nu(\theta)\,\lambda_R(B)\,TC\left(h_\theta, B; \theta'\right) \end{aligned} \tag{B.23}$$

where $\theta' = \max\left\{\widetilde{\theta} \in \Theta : \widetilde{\theta} < \theta\right\}$ as before, and we have again made use of the restriction

$$\sum_{\theta\in\Theta} \pi_\Theta(\theta|h_\theta)\,\nu(\theta) = 0$$

in obtaining this expression. Since the left-hand side of (B.23) is now known to be zero, the result follows from trivial manipulation.

B.3 Derivation of equation 5.15

For this class of ν vectors we will still have condition (B.21), but the second sum on the right-hand side no longer collapses to zero. Instead, using (B.13) we can write:

$$\begin{aligned} & \sum_{\widetilde{\theta}\in\Theta} \pi_\Theta\left(\widetilde{\theta}|\left(h_\theta,\theta\right)\right)\nu\left(\widetilde{\theta}\right)\lambda_V^0\left(\left(h_\theta,\theta\right),B;\widetilde{\theta}\right) \\ & + \sum_{\widetilde{\theta}\in\Theta\backslash\underline{\theta}} \pi_\Theta\left(\theta'|\left(h_\theta,\theta\right)\right)\nu\left(\widetilde{\theta}\right)\lambda_{\widehat{V}}^0\left(\left(h_\theta,\theta\right),B;\widetilde{\theta}\right) \\ = & \ \beta\sum_{\widetilde{\theta}\in\Theta} \pi_\Theta\left(\widetilde{\theta}|\left(h_\theta,\theta''\right)\right)\nu\left(\widetilde{\theta}\right)\lambda_{\widehat{V}}^1\left(\left(h_\theta,\theta\right),B;\widetilde{\theta}\right) \\ = & \ \beta\lambda_{\widehat{V}}^0\left(h_\theta,B;\theta''\right) \end{aligned} \tag{B.24}$$

where $\theta' = \max\left\{\widetilde{\theta}\in\Theta : \widetilde{\theta} < \theta\right\}$ and $\theta'' = \min\left\{\widetilde{\theta}\in\Theta : \widetilde{\theta} > \theta\right\}$, and we make use of the fact optimal 'backward-looking' policy implies $B\left(B\right) = B$ to simplify the notation slightly.

We have from (B.19):

$$\lambda_{\widehat{V}}^0\left(h_\theta,B;\theta''\right) = \lambda_R\left(B\right)TC\left(h_\theta,B;\theta\right) \tag{B.25}$$

and so proceeding in the same manner as for equation (B.23) we have:

$$\begin{aligned} & \beta TC\left(h_\theta,B;\theta\right) \\ = & -\sum_{\widetilde{\theta}\in\Theta} \pi_\Theta\left(\widetilde{\theta}|\left(h_\theta,\theta\right)\right)\nu\left(\widetilde{\theta}\right)\left[MC\left(\left(h_\theta,\theta\right),B;\widetilde{\theta}\right) + TC\left(\left(h_\theta,\theta\right),B;\widetilde{\theta}\right)\right] \\ & + \sum_{\widetilde{\theta}\in\Theta\backslash\underline{\theta}} \pi_\Theta\left(\theta'|h_\theta\right)\nu\left(\widetilde{\theta}\right)TC\left(\left(h_\theta,\theta\right),B;\theta'\right) \end{aligned} \tag{B.26}$$

for $\theta' = \max\left\{\theta\in\Theta : \theta < \widetilde{\theta}\right\}$. The result follows.

C References

References

[1] Abreu, D., D. Pearce and E. Stacchetti (1990), 'Toward a Theory of Discounted Repeated Games with Imperfect Monitoring', *Econometrica*, 58, 1041–63.

[2] Alesina, A. and L.H. Summers (1993), 'Central Bank Independence and Macroeconomic Performance: Some Comparative Evidence', *Journal of Money, Credit and Banking*, 25, 151–162.

[3] Atkeson, A. and R.E. Lucas Jr. (1992), 'On Efficient Distribution with Private Information', *Review of Economic Studies*, 59, 427–453.

[4] Barro, R.J. and D.B. Gordon (1983), 'A positive theory of monetary policy in a natural-rate model', *Journal of Political Economy*, 91, 589–610.

[5] Blake, A.P. (2001), 'A "Timeless Perspective" on Optimality in Forward-Looking Rational Expectations Models', NIESR Discussion Papers 188, National Institute of Economic and Social Research.

[6] Blake, A.P. and T. Kirsanova (2008), 'Discretionary Policy and Multiple Equilibria in LQ RE Models', Working Paper, University of Exeter (available at http://ssrn.com/abstract=943032).

[7] Brendon, C.F. (2011), 'Applying perturbation methods to dynamic tax problems with a continuum of types', Manuscript (available on request).

[8] Broer, T., M. Kapička and J. Klein (2011), 'Consumption Risk Sharing under Private Information when Earnings are Persistent', Manuscript (available at http://paulklein.se/newsite/research/risksharing.pdf).

[9] Currie, D. and P. Levine (1985), 'Simple Macropolicy Rules for the Open Economy', *Economic Journal*, 95, 60–70.

[10] Calvo, G. (1983), 'Staggered prices in a utility-maximizing framework', *Journal of Monetary Economics*, 12, 383–98.

[11] da Costa, C., and I. Werning (2002), 'Commodity Taxation and Social Insurance', Manuscript (available at http://economics.mit.edu/files/1265).

[12] Damjanovic, T., V. Damjanovic, and C. Nolan (2008), 'Unconditionally optimal monetary policy', *Journal of Monetary Economics*, 55, 491–500.

[13] Dasgupta, P. (1982), 'Utilitarianism, Information and Rights', in A. Sen and B. Williams (eds.), *Utilitarianism and Beyond*, Cambridge, Cambridge University Press.

[14] Diamond, P.A., and J.A. Mirrlees (1978), 'A Model of Social Insurance with Variable Retirement', *Journal of Public Economics*, 10, 295–336.

[15] Diamond, P.A. and E. Saez (2011), 'The Case for a Progressive Tax: From Basic Research to Policy Recommendations', *Journal of Economic Perspectives*, 25, 165–190.

[16] Farhi, E. and I. Werning (2007), 'Inequality and Social Discounting', *Journal of Political Economy*, 115, 365–402.

[17] Farhi, E. and I. Werning (2010), 'Insurance and Taxation over the Life Cycle', Manuscript (available at http://econ-www.mit.edu/files/6588).

[18] Fernandes, A. and C. Phelan (2000), 'A Recursive Formulation for Repeated Agency with History Dependence', *Journal of Economic Theory*, 91, 223–247.

[19] Giannoni, M. P., and M. Woodford (2002), 'Optimal Interest-Rate Rules 1: General Theory', NBER Working Paper No. 9419.

[20] Golosov, M., N. Kocherlakota and A. Tsyvinski (2003), 'Optimal Indirect and Capital Taxation', *Review of Economic Studies*, 70, 569–587.

[21] Golosov, M., M. Troshkin and A. Tsyvinski (2011), 'Optimal Taxation: Merging Micro and Macro Approaches', *Journal of Money, Credit and Banking*, 43, 147–174.

[22] Golosov, M., A. Tsyvinski and I. Werning (2006), 'New Dynamic Public Finance: A User's Guide', in *NBER Macroeconomic Annual 2006*, Cambridge MA, MIT Press.

[23] Greenwald, B., and J.E. Stiglitz (1986), 'Externalities in Economies with Imperfect Information and Incomplete Markets', *Quarterly Journal of Economics*, 101, 229–264.

[24] Ireland, P. N. (1997), 'Sustainable Monetary Policies', *Journal of Economic Dynamics and Control*, 22, 87–108.

[25] M. Kapička (2010), 'Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First-Order Approach', Manuscript (available at http://www.econ.ucsb.edu/~mkapicka/persistent.pdf), and forthcoming in *Review of Economic Studies*.

[26] Kocherlakota, N.R. (1996), 'Implications of Efficient Risk Sharing without Commitment', *Review of Economic Studies*, 63, 595-609.

[27] Kocherlakota, N.R. (2005), 'Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation', *Econometrica*, 73, 1587–1621.

[28] Kocherlakota, N.R. (2011), *The New Dynamic Public Finance*, Princeton, Princeton University Press.

[29] Kydland, F.E. and E.C. Prescott (1977), 'Rules rather than discretion: the inconsistency of optimal plans', *Journal of Political Economy*, 85, 473–491.

[30] Ljungqvist, L. and T. J. Sargent (2004), *Recursive Macroeconomic Theory*, MIT Press

[31] Lucas, R. Jr. and N.L. Stokey (1983), 'Optimal fiscal and monetary policy in an economy without capital', *Journal of Monetary Economics*, 12, 55–93.

[32] Mankiw, N.G., M. Weinzierl and D. Yagan (2009), 'Optimal Taxation in Theory and Practice' *Journal of Economic Perspectives*, 23, 147–174.

[33] Marcet, A. and R. Marimon (1992), 'Communication, Commitment and Growth', *Journal of Economic Theory*, 58, 219–249

[34] Marcet, A. and R. Marimon (1998), 'Recursive Contracts', Economics Working Papers eco98/37, European University Institute.

[35] Marcet, A. and R. Marimon (2011), 'Recursive Contracts', CEP Discussion Paper 1055, Centre for Economic Performance, LSE.

[36] Mele, A. (2011), 'Repeated Moral Hazard and Recursive Lagrangeans', Working Paper (available at http://sites.google.com/site/meleantonio/#research).

[37] Messner, M. and N. Pavoni (2004), 'On the Recursive Saddle Point Method', Working Paper (available at ftp://ftp.igier.uni-bocconi.it/wp/2004/255.pdf).

[38] Messner, M., N. Pavoni and C.M. Sleet (2011), 'Recursive Methods for Incentive Problems', Manuscript (available at SSRN: http://ssrn.com/abstract=1792816).

[39] Mirrlees, J.A. (1971), 'An Exploration in the Theory of Optimum Income Taxation', *Review of Economic Studies*, 38, 175–208.

[40] Oudiz, G. and J. Sachs (1985), 'International policy coordination in dynamic macroeconomic models'; in W. H. Buiter and R. C. Marston (Eds.), *International Economic Policy Coordination*, Cambridge: Cambridge University Press.

[41] Pavan, A., I. Segal and J. Toikka (2011), 'Dynamic Mechanism Design: Incentive Compatibility, Profit Maximization and Information Disclosure', Manuscript (available at http://www.stanford.edu/~isegal/dmd.pdf).

[42] Phelan, C. (2006), 'Opportunity and Social Mobility', *Review of Economic Studies*, 73, 487–505.

[43] Ramsey, F.P. (1928), 'A Mathematical Theory of Saving', *The Economic Journal*, 38, 543–559.

[44] Rawls, J. (1971), *A Theory of Justice*, Oxford: Clarendon Press.

[45] Roberts, K. (2000), 'A Reconsideration of the Optimal Income Tax', in Hammond, P.J., and G.D. Myles (eds.), *Incentives and Organization: Papers in Honour of Sir James Mirrlees* (Oxford: Oxford University Press).

[46] Rogoff, K. (1985), 'The Optimal Degree of Commitment to an Intermediate Monetary Target', *Quarterly Journal of Economics*, 100, 1169–1190.

[47] Saez, E. (2001), 'Using Elasticities to Derive Optimal Income Tax Rates', *Review of Economic Studies*, 68, 205–229.

[48] Sleet, C.M. and Ş. Yeltekin (2010a), 'The Recursive Lagrangian Method: Discrete Time', Manuscript (available at http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.167.7942).

[49] Sleet, C.M. and Ş. Yeltekin (2010b), 'Misery and Luxury: Long Run Outcomes with Private Information', Manuscript (available at https://student-3k.tepper.cmu.edu/gsiadoc/WP/2011-E19.pdf).

[50] Svensson, L.E.O. (1999), 'How Should Monetary Policy Be Conducted In An Era Of Price Stability?', CEPR Discussion Paper 2342.

[51] Taylor, J. B. (1979), 'Estimation and Control of a Macroeconomic Model with Rational Expectations', *Econometrica*, 47, 1267–1286.

[52] Thomas, J.P. and T. Worrall (1990), 'Income Fluctuation and Asymmetric Information: An Example of a Repeated Principal-Agent Model', *Journal of Economic Theory*, 51, 367–390.

[53] Walsh, C.E. (1995), 'Optimal contracts for central bankers', *American Economic Review*, 85,150–167.

[54] Woodford, M. (2003), *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton: Princeton University Press.

[55] Woodford, M. (2010), 'Optimal Monetary Stabilization Policy'; forthcoming in B.M. Friedman and M. Woodford (Eds.), *Handbook of Monetary Economics, Volume 3B,* Amsterdam: Elsevier.