Applying perturbation analysis to dynamic optimal tax problems*

Charles Brendon
University of Oxford

April 13, 2012

Abstract

This paper shows how to derive a complete set of optimality conditions characterising the solution to a dynamic optimal income tax problem in the spirit of Mirrlees (1971), under the assumption that a 'first-order' approach to incentive compatibility is valid. The method relies on constructing a novel class of perturbations to the consumption-output allocations of agents, in a manner that preserves all incentive compatibility constraints that are binding at a putative optimum. 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, to characterise the dynamic evolution of the effective labour tax rate, and to prove a number of other novel results about optimal tax wedges.

JEL Classification Codes: D82, E61, H21, H24
Keywords: New Dynamic Public Finance, first-order approach, non-separable preferences, inverse Euler condition

*I would like to thank Mikhail Golosov, Albert Marcet, Kevin Roberts, Rick van der Ploeg, Simon Wren-Lewis and, in particular, Antonio Mele for providing helpful comments relating to this work. All errors are mine.

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 paper 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 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 system of simultaneous equations (with no need to use dynamic programming techniques). More significantly, the conditions that we derive imply several novel qualitative 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 supply levels are non-separable in utility. 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 does not generalise to the empirically relevant case in which they are complements. Similarly, and again using our generalisation of the inverse Euler condition, we show that the long-run 'immiseration' results that can 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

[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).

responsible for immiseration.

In addition to savings distortions, another important economic variable that features in dynamic optimal tax analyses is the implicit labour tax wedge – that is, the wedge between an agent's within-period marginal rate of substitution between consumption and production, and the marginal rate of transformation. In an important recent paper Farhi and Werning (2011) provide a novel characterisation of the dynamics of this wedge when the first-order approach is valid, focusing particular attention on the case of separable preferences when labour is supplied isoelastically. In this paper we provide an analogous expression that holds under general preference specifications, and show how it can be interpreted in terms of a dynamic 'efficiency-equity trade-off' of the kind often emphasised in the traditional public finance literature.[3]

Accounting for these results on a technical level, 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.[4] But 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. The dynamic optimality condition for the labour wedge is precisely this condition – implying that any persistence in the labour wedge must be *inherited* from persistence in the productivity process, rather than being induced directly by the optimal plan.

We are additionally able to confirm some more common observations regarding within-period character of the labour wedge: that this should always be weakly positive (it is never optimal to subsidise work), and zero at any upper support for the productivity distribution. Once again, these familiar properties of the static model generalise with remarkable ease to dynamic settings.

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

[3] Farhi and Werning do provide an expression for the evolution of the labour wedge under general preference specifications, but its dependence on multipliers makes this comparatively difficult to interpret.

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

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).[5]

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 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 (2011a) in a dynamic Mirrlees model; 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.

[5] 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).

An important refinement to this method – particularly in the context of the present paper – has been provieded by Kapička (2011). This extends the general work of Pavan, Segal and Toikka (2011), and illustrates 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 (2011) and Golosov, Troshkin and Tsyvinski (2011b).

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. Alongside important work by Mele (2011), extending the work of Marcet and Marimon to repeated hidden action problems, this method has recently been applied to optimal dynamic tax policy by Sleet and Yeltekin (2010b). 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: this is the so-called 'inverse Euler condition', 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.

In this paper we show that the perturbation logic behind the inverse Euler condition generalises considerably, provided one is willing to accept the first-order approach to incentive compatibility. This remains something of a leap. The only known general conditions under which the validity of the first-order approach can be assured impose requirements on the properties of solutions obtained under it;[6] these conditions cannot, therefore, be confirmed *ex-ante*, based on the model's priors alone. But as noted above, a number of papers have made good progress understanding dynamic tax models by applying computational techniques that assume the first-order approach *is* valid, confirming this *ex-post*. The perturbation method presented here is intended complement these analyses, retaining their caveats.

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:

$$U_t = E_t \sum_{s=0}^{\infty} \beta^s u\left(c_{t+s}, y_{t+s}; \theta_{t+s}\right) \tag{1}$$

where $u : \mathbb{R}_+^3 \to \mathbb{R}$. 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.[7] 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

[6] See, for instance, Theorem 5 in Kapička (2011).

[7] 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.

paper (see Brendon (2011)). We also assume an infinite horizon, though none of the optimality conditions that we derive is dependent on this perspective.[8]

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:

$$u_{y\theta} - u_{yy}\frac{u_\theta}{u_y} > 0 \tag{2}$$

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

$$u_{c\theta} - u_{cy}\frac{u_\theta}{u_y} = 0 \tag{3}$$

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:

$$u\left(c'', y''; \theta''\right) - u\left(c', y'; \theta''\right) > u\left(c'', y''; \theta'\right) - u\left(c', y'; \theta'\right) \tag{4}$$

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

[8] 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.

curve in output-consumption space ($\frac{dc}{dy} = -\frac{u_y}{u_c}$); although (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 (4) implies for all (c, y) pairs:

$$-\frac{u_y\left(c, y; \theta''\right)}{u_c\left(c, y; \theta''\right)} < -\frac{u_y\left(c, y; \theta'\right)}{u_c\left(c, y; \theta'\right)} \tag{5}$$

This follows directly from the fact that indifference curves in consumption-output space must be 'flattening' as θ increases, provided (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,[9] 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.[10] 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

[9] 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.

[10] 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.

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.

The policymaker's primal choice problem is, therefore:

$$\max_{\{c_t(\theta^\infty),y_t(\theta^\infty)\}_{t=1}^{\infty}} \int_{\Theta^\infty} \sum_{t=1}^{\infty} \beta^{t-1} u\left(c_t\left(\theta^\infty\right), y_t\left(\theta^\infty\right); \theta_t\right) d\pi_\Theta\left(\theta^\infty\right) \tag{6}$$

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

$$\begin{aligned} & \int_{\Theta^\infty} \sum_{s=0}^{\infty} \beta^s u\left(c_{t+s}\left(\theta^\infty\right), y_{t+s}\left(\theta^\infty\right); \theta_{t+s}\right) d\pi_\Theta\left(\theta^\infty | \theta^t\right) \\ \geq & \int_{\Theta^\infty} \sum_{s=0}^{\infty} \beta^s u\left(c_{t+s}\left(\widehat{\theta}^\infty\left(\theta^\infty\right)\right), y_{t+s}\left(\widehat{\theta}^\infty\left(\theta^\infty\right)\right); \theta_{t+s}\right) d\pi_\Theta\left(\theta^\infty | \theta^t\right) \end{aligned} \tag{7}$$

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:

$$\int_{\Theta^\infty} \left[c_t\left(\theta^\infty\right) - y_t\left(\theta^\infty\right)\right] d\pi_\Theta\left(\theta^\infty\right) + A_{t+1} = R_t A_t \tag{8}$$

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 $\lim_{s\to\infty}\left[\left(\prod_{r=1}^{s} R_{t+r}^{-1}\right) A_{t+s}\right] = 0$.[11]

3 Full information benchmark

In a manner equivalent to Kapička (2011) 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 $\theta_t' = \max\{\theta \in \Theta : \theta < \theta_t\}$. 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

[11] 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.

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 (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 $i \in [0,1]$:

$$u_c\left(c_t^i, y_t^i; \theta_t^i\right) = -u_y\left(c_t^i, y_t^i; \theta_t^i\right) \tag{9}$$

$$u_c\left(c_t^i, y_t^i; \theta_t^i\right) = \beta R_{t+1} \sum_{\theta_{t+1}^i \in \Theta} u_c\left(c_{t+1}^i, y_{t+1}^i; \theta_{t+1}^i\right) \pi_\Theta\left(\theta_{t+1}^i | \theta_t^i\right) \tag{10}$$

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

$$u_c\left(c_t^i, y_t^i; \theta_t^i\right) = u_c\left(c_t^j, y_t^j; \theta_t^j\right) \tag{11}$$

   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. 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.

For all periods $t \geq 1$, define $\widehat{\theta}^{\infty}_{m,t} : \Theta^{\infty} \rightarrow \Theta^{\infty}$ 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:

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

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 $\widehat{\theta}^{\infty}_{m,t}\left(\theta^{\infty}\right) = \theta^{\infty}$. If incentive compatabilty is said to be holding 'downwards', the following is true:

$$\begin{aligned} & \int_{\Theta^{\infty}} \sum_{s=0}^{\infty} \beta^s u\left(c_{t+s}\left(\theta^{\infty}\right), y_{t+s}\left(\theta^{\infty}\right); \theta_{t+s}\right) d\pi_{\Theta}\left(\theta^{\infty}|\theta^t\right) \qquad (12) \\ \geq & \int_{\Theta^{\infty}} \sum_{s=0}^{\infty} \beta^s u\left(c_{t+s}\left(\widehat{\theta}^{\infty}_{m,t}\left(\theta^{\infty}\right)\right), y_{t+s}\left(\widehat{\theta}^{\infty}_{m,t}\left(\theta^{\infty}\right)\right); \theta_{t+s}\right) d\pi_{\Theta}\left(\theta^{\infty}|\theta^t\right) \end{aligned}$$

So the agent with history θ^t is just indifferent between reporting θ_t truthfully and 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 (12) 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, $\widehat{\theta}^{\infty}_{a,t} : \Theta^{\infty} \rightarrow \Theta^{\infty}$, defined by:

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

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 equation (12) implies:

$$\int_{\Theta^\infty} \sum_{s=0}^{\infty} \beta^s u\left(c_{t+s}\left(\theta^\infty\right), y_{t+s}\left(\theta^\infty\right); \theta_{t+s}\right) d\pi_\Theta\left(\theta^\infty | \theta^t\right) \tag{13}$$

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

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 $W : \Theta \times \Theta \times \Theta^{t-1} \rightarrow \mathbb{R}$ 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 (13) 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:

$$W\left(\theta_t; \theta_t, \widehat{\theta}^{t-1}\right) \geq W\left(\theta_t'; \theta_t, \widehat{\theta}^{t-1}\right) \tag{14}$$

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 model, incentive compatibility must be imposed 'downwards'. It should be stressed that in general condition (14) 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 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

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

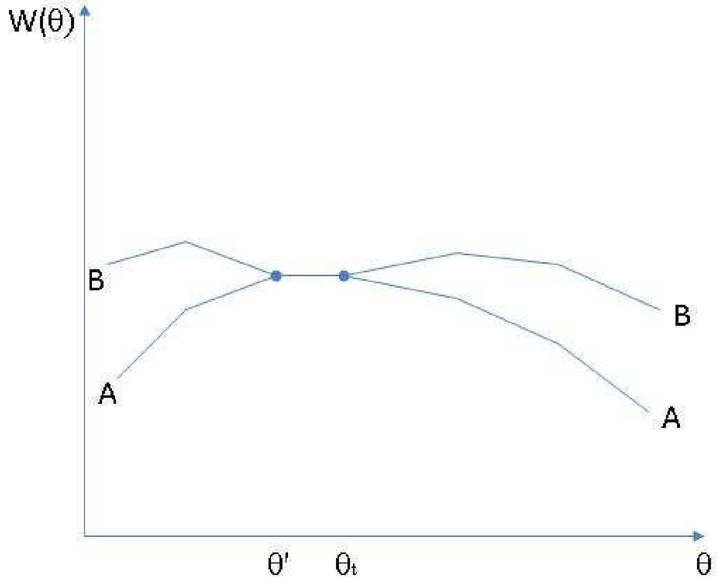

Figure 1: Local incentive compatibility need not imply global.

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 (14) 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 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 (14) 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 (14) 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* $\theta'_t=\max\left\{\theta\in\Theta:\theta<\theta_t\right\}$ *and* $\underline{\theta}=\min\left\{\theta\in\Theta\right\}$*).*

This is a natural translation to our discrete-type setting of Theorem 5 in Kapička (2011).[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 (14) 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 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 (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 (14) 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. ■

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

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 (14).

In what follows we refer to the problem of policy choice from among the set of direct revelation mechanisms satisfying condition (14) 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 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. We therefore make the following assumption throughout:

Assumption 1: The solution to the relaxed problem also solves the general problem.

Since we are interested in perturbations about the optimum it also helps to assume:

Assumption 2: The solution to the relaxed problem is interior, in the sense that $c_t(\theta^\infty) > 0$ and $y_t(\theta^\infty) > 0$ for all $t \geq 1$ and all $\theta^\infty \in \Theta^\infty$.

We justify this by appeal to the Inada conditions that we have imposed. The results that we have should generalise to corner solutions, but only with substantial additional notational baggage.

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

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

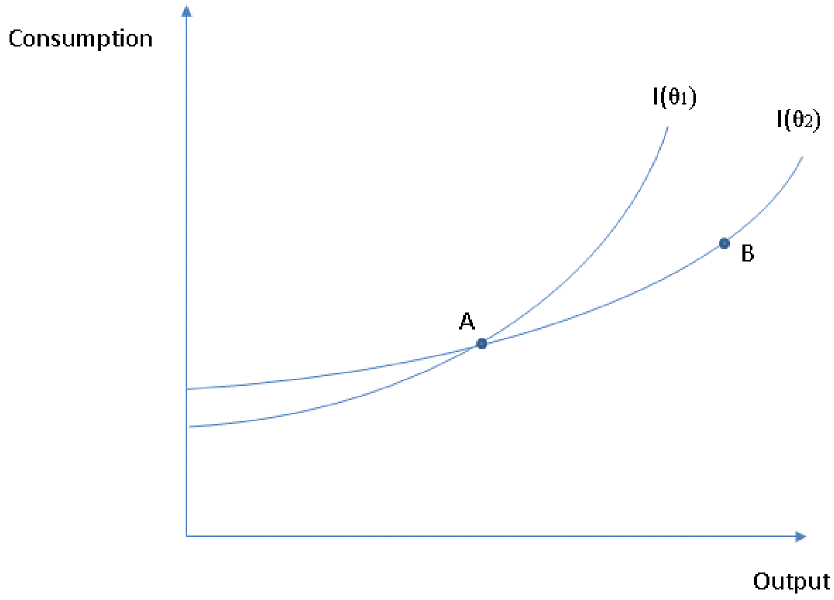

Figure 2: Graphical introduction to the policy trade-off

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 2 shows such a mapping.

The two within-period indifference curves are drawn for arbitrary distinct types θ_1 and θ_2, with $\theta_2 > \theta_1$. 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 2 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 (14) replaces the complete constraint set, and note again ours twin assumptions that the solution to this problem coincides with the solution to the general problem in which the full constraint set is imposed, and that consumption and output allocations are strictly positive for all agents at all finite horizons. 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 $\widehat{\theta}_t \in \Theta$,[15] 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 $\Delta : \mathbb{R} \to \mathbb{R}^{2N}$) for some relevant parameter δ (perhaps the consumption or utility increment implied by the given

[15] 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.

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 $\delta = 0$.

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.[16] 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

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

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*.

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.[17] 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 $u_c = -u_y$.[18] The logic generalises to the intertemporal model, as the following makes clear.

Proposition 2 ***No distortion at the top:*** *In all time periods $t \geq 1$ and (if $t > 1$) for all past reporting 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 $(0, 0)$

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

[18] 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.

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:

$$u\left(c_t^* + \delta, y_t^* + f(\delta); \theta_t'\right) = u\left(c_t^*, y_t^*; \theta_t'\right) \tag{15}$$

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:

$$\pi_\Theta\left(\theta_t'|\theta^{t-1}\right)\pi_\Theta\left(\theta^{t-1}\right)[f'(0) - 1] = 0 \tag{16}$$

Probabilities are non-zero, so this implies:

$$f'(0) = 1 \tag{17}$$

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

$$u_c\left(c_t^*, y_t^*; \theta_t'\right) + u_y\left(c_t^*, y_t^*; \theta_t'\right) f'(0) = 0 \tag{18}$$

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.[19] 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

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

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 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 $\alpha : \mathbb{R}^2_+ \times \Theta \to \mathbb{R}$, as follows:

$$\alpha(c, y; \theta) = \frac{u_c(c, y; \theta) - u_c(c, y; \theta')}{u_y(c, y; \theta') - u_y(c, y; \theta)} \tag{19}$$

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 $\alpha(c, y; \theta) = 0$.[20]

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.[21] 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, 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 $\alpha > 0$.[22] That is, higher production must accompany higher

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

[21] 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.

[22] 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 (3) further implies $u_{c\theta} < 0$ for Edgeworth complements and $u_{c\theta} > 0$ for Edgeworth substitutes.

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 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:*** *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:*

$$
\begin{aligned}
& R_{t+1}\beta \frac{1-\alpha\left(c_t^*, y_t^*; \theta_t\right)}{u_c\left(c_t^*, y_t^*; \theta_t\right)+u_y\left(c_t^*, y_t^*; \theta_t\right)\alpha\left(c_t^*, y_t^*; \theta_t\right)} \qquad (20) \\
& = \sum_{\theta_{t+1}\in\Theta} \pi_\Theta\left(\theta_{t+1}|\theta^t\right) \frac{1-\alpha\left(c_{t+1}^*, y_{t+1}^*; \theta_{t+1}\right)}{u_c\left(c_{t+1}^*, y_{t+1}^*; \theta_{t+1}\right)+u_y\left(c_{t+1}^*, y_{t+1}^*; \theta_{t+1}\right)\alpha\left(c_{t+1}^*, y_{t+1}^*; \theta_{t+1}\right)}
\end{aligned}
$$

where c_{t+1}^* and y_{t+1}^* are given by the relevant entries in $X_{t+1}^*\left(\theta^t\right)$.[23]

A full proof is given in the appendix. Heuristically, 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

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

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 (20) 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 (20) 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.

The marginal cost term that features in (20) will appear repeatedly in the analysis that follows. To avoid repeatedly writing out a relatively unwieldy object, we refer to it as $MC(c, y; \theta)$:

$$MC(c, y; \theta) \equiv \frac{1 - \alpha(c, y; \theta)}{u_c(c, y; \theta) + u_y(c, y; \theta)\,\alpha(c, y; \theta)} \tag{21}$$

5.4.3 Implications for optimal savings wedges

This generalised inverse Euler condition provides useful theoretical insight, but our main aim is to exploit it to make qualitative statements about the character of optimal tax distortions. The obvious area to turn to is optimal savings wedges – that is, how the policymaker might wish to drive a wedge between agents' intertemporal marginal rates of substituion and the real interest rate. The first point to note here is that (20) provides qualified analytical support to the numerical result of Golosov, Tsyvinski and Werning (2006) that the optimal savings wedge to insert in the consumption Euler equation *could* be negative for some agents, at least in the weak sense that under some preference structures we cannot analytically rule *out* the inequality:

$$u_c(c_t^*, y_t^*; \theta_t) > \beta R_{t+1} \sum_{\theta_{t+1} \in \Theta} \pi_\Theta\left(\theta_{t+1} | \theta^t\right) u_c\left(c_{t+1}^*, y_{t+1}^*; \theta_{t+1}\right) \tag{22}$$

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 (22) 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:

$$\sum_{\theta_{t+1} \in \Theta} \pi_\Theta\left(\theta_{t+1} | \theta^t\right) \left[\frac{1}{u_c\left(c_{t+1}^*, y_{t+1}^*; \theta_{t+1}\right)}\right] \geq \left[\sum_{\theta_{t+1} \in \Theta} \pi_\Theta\left(\theta_{t+1} | \theta^t\right) u_c\left(c_{t+1}^*, y_{t+1}^*; \theta_{t+1}\right)\right]^{-1} \tag{23}$$

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 (20) in the case of separable preferences ($\alpha(c, y; \theta) = 0$) confirms that savings are indeed deterred:[24]

$$u_c(c_t^*, y_t^*; \theta_t) \leq \beta R_{t+1} \sum_{\theta_{t+1} \in \Theta} \pi_\Theta\left(\theta_{t+1} | \theta^t\right) u_c\left(c_{t+1}^*, y_{t+1}^*; \theta_{t+1}\right) \tag{24}$$

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

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

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 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.[25] 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 (20), 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 (24) in the non-separable case is only possible in very particular circumstances. Fortunately there is a natural generalisation of the 'deterred savings' inequality that will apply more widely; but first we present the results that we can state regarding this standard consumption Euler condition.

[25] 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).

Proposition 4 ***Deterred savings (1):*** *Suppose 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 (24) 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 $\theta_t = \max\{\theta \in \Theta\}$.*

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 7 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 (24). 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:

$$u_y(\theta_t) = \beta R_{t+1} \sum_{\theta_{t+1} \in \Theta} \pi_\Theta\left(\theta_{t+1} | \theta^t\right) u_y(\theta_{t+1}) \tag{25}$$

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:

$$\begin{aligned} & \frac{u_c(\theta_t) + u_y(\theta_t)\,\alpha(\theta_t)}{1 - \alpha(\theta_t)} \qquad (26) \\ = \; & \beta R_{t+1} \sum_{\theta_{t+1} \in \Theta} \pi_\Theta\left(\theta_{t+1} | \theta^t\right) \frac{u_c(\theta_{t+1}) + u_y(\theta_{t+1})\,\alpha(\theta_{t+1})}{1 - \alpha(\theta_{t+1})} \end{aligned}$$

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 (26) 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 (26) does, is of equal validity to the consumption Euler equation in characterising optimal dynamic behaviour under autarky.

The useful feature of equation (26) 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):*** *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:*

$$\begin{aligned} & \frac{u_c\left(\theta_t\right)+u_y\left(\theta_t\right)\alpha\left(\theta_t\right)}{1-\alpha\left(\theta_t\right)} \\ \leq \quad & \beta R_{t+1}\sum_{\theta_{t+1}\in\Theta}\pi_\Theta\left(\theta_{t+1}|\theta^t\right)\frac{u_c\left(\theta_{t+1}\right)+u_y\left(\theta_{t+1}\right)\alpha\left(\theta_{t+1}\right)}{1-\alpha\left(\theta_{t+1}\right)} \end{aligned} \tag{27}$$

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 $\alpha\left(\theta_t\right)<0$, so for the preferences we are focusing on we must always have:

$$\frac{u_c\left(\theta_t\right)+u_y\left(\theta_t\right)\alpha\left(\theta_t\right)}{1-\alpha\left(\theta_t\right)}>0 \tag{28}$$

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

$$\begin{aligned} & \sum_{\theta_{t+1}\in\Theta}\pi_\Theta\left(\theta_{t+1}|\theta^t\right)\left[\frac{u_c\left(\theta_{t+1}\right)+u_y\left(\theta_{t+1}\right)\alpha\left(\theta_{t+1}\right)}{1-\alpha\left(\theta_{t+1}\right)}\right]^{-1} \\ \geq \quad & \left[\sum_{\theta_{t+1}\in\Theta}\pi_\Theta\left(\theta_{t+1}|\theta^t\right)\frac{u_c\left(\theta_{t+1}\right)+u_y\left(\theta_{t+1}\right)\alpha\left(\theta_{t+1}\right)}{1-\alpha\left(\theta_{t+1}\right)}\right]^{-1} \end{aligned} \tag{29}$$

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 (29) is also the right-hand side of equation (20); the inequality in the Proposition then follows straightforwardly from using (20) in (29). ■

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 5 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 (20) 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 the paper.

5.5 General intratemporal 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 appled (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 how to construct an additional set of $N-1$ 'intratemporal' perturbations that can be applied in the iid case – changing within-period outcomes across agents with a common prior reporting history, but leaving allocations in all other time periods unaffected. The optimality conditions associated with these perturbations will be enough to close the model, in the sense that they provide as many restrictions as there degrees of freedom left to the policymaker. Because they are conditions that must apply irrespective of the dynamic setting of the model, they also provide novel insight into the 'static' optimal income tax model, of the type originally set out in Mirrlees (1971).

The intuition that we exploit is the following. Suppose one wished to perturb the within-period *utility* levels at time t of all (truthful) agents with prior reporting history $\widehat{\theta}^{t-1}$, in proportions corresponding to the elements of some $N \times 1$ vector ν – that is, increasing the utility of type θ_t^1 by an amount $\nu_1 \delta$, of type θ_t^2 by an amount $\nu_2 \delta$, and so on, for some common scalar δ. If one is to remain within the constraint set of the relaxed problem, these perturbations must be engineered in such a way that the change in utility received by the (truth-telling) nth agent is equal to the change in the utility that this agent could obtain by mimicking those of type θ_t^{n-1}. This in turn can be ensured by delivering the perturbation to θ_t^{n-1}'s utility along an appropriately-chosen dimension in consumption-output space, exploiting the fact that along all dimensions *apart* from those highlighted in presenting our generalised inverse Euler condition, utility changes to mimickers *differ* from utility changes to truth-tellers.

Collectively these utility changes at t could potentially affect incentive compatibility at $t-1$ – that is, the incentives to report $\widehat{\theta}^{t-1}$. This will be the case whenever they increase or decrease the expected value of within-period utility across types at t. To prevent this we make sure that the perturbations at t are constructed so as not change within-period utility at all in expectation. This prevents any need to perturb prior allocations, and thus keeps the focus restricted to a single time period – yielding 'intratemporal' optimality conditions as required. Note that the iid assumption will ensure that both truth-tellers and (one-higher) mimickers at $t-1$ experience the same *ex-ante* change to their expected within-period utility from reporting $\widehat{\theta}^{t-1}$, even though the proposed perturbation has differential effects across types at t.[26]

Figure 3 illustrates diagramatically the type of construction we have in mind. Suppose that under the chosen ν vector $\nu_n = 0$ and $\nu_{n+1} = 1$. This means that for any given δ we wish to increase the within-period utility of the $n+1$th agent by an amount δ, but leave unchanged the utility of the nth. The latter will be ensured only if the allocation to the nth agent is perturbed along that agent's within-period indifference curve. To preserve (exact) incentive compatibility for the $n+1$th agent, this movement must deliver δ units of extra utility to that agent when mimicking. The perturbation to the $n+1$th agent's allocation will likewise increase that agent's truth-telling utility by an amound δ, and should be chosen based on the required increment to mimicking agents of type θ^{n+1} (omitted from the diagram).

Since we are free to select in this way *any* utility increment vector ν whose expected value across types is zero *ex ante*, the associated class of perturbations is potentially very large.

[26] 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 $t-1$.

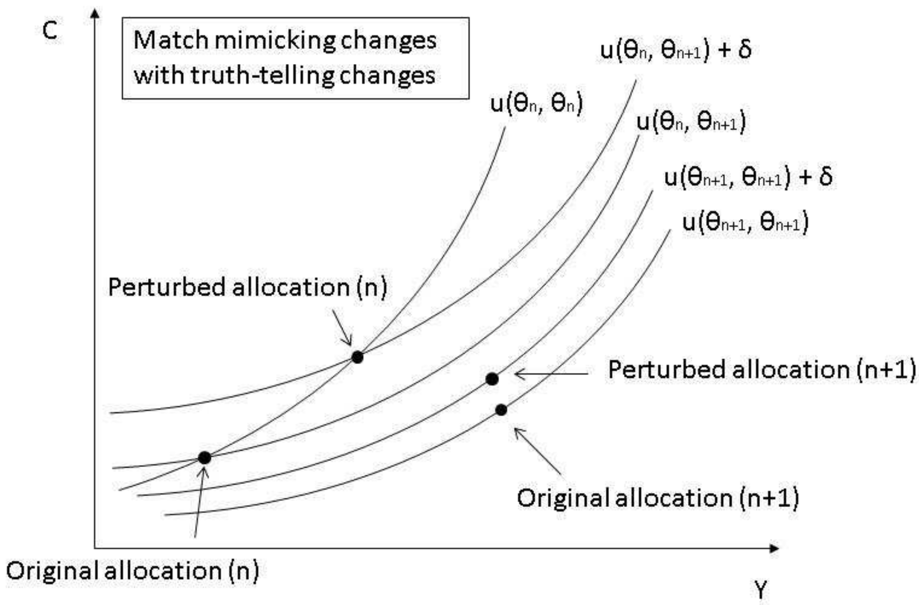

Figure 3: Engineering incentive-compatible utility changes

5.5.2 Analytical treatment

We now present the arguments more formally. It is useful first 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 is 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:

$$u_c(c, y, \theta)(1 - \tau(c, y, \theta)) = -u_y(c, y, \theta) \tag{30}$$

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.

A first step towards obtaining the general optimality conditions we are after is the following Lemma, the proof of which is in the appendix:

Lemma 1 *Suppose that type draws are iid across agents and time. Fix a vector $\nu \in \mathbb{R}^N$ (whose nth element is denoted ν_n) such that:*

$$\sum_{n=1}^{N} \pi_\Theta(\theta^n)\, \nu_n = 0$$

For all time periods $t \geq 1$ and any prior reporting history $\widehat{\theta}^{t-1} \in \Theta^{t-1}$ it is possible to perturb the set of optimal allocations $X_t^*\left(\widehat{\theta}^t\right)$ 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^n by an amount $\nu_n\delta$ in period t, for any scalar δ satisfying $|\delta| < \varepsilon$ for some $\varepsilon > 0$.

This confirms that any incremental vector that delivers zero *expected* utility across types from the perspective of previous time periods can be engineered through changes to within-period allocations alone, without affecting incentive compatibility constraints at any horizon. The proof is just a formal statement of the intuition developed graphically in the previous time period.

What really matters if we want to say something about optimality, though, is the impact that perturbations of this kind have on resources at the margin as the scalar δ is moved away from zero. If we start at an optimum, this marginal cost must itself be zero. It turns out that an important object when assessing the overall marginal resource cost of any perturbation is the specific marginal cost to the policymaker of a movement along the nth agent's utility curve, by an amount just sufficient to reduce by one unit the utility that could be obtained by the $n+1$th agent when mimicking the nth. In terms of the graphical analysis just presented, this is a 'downwards' movement along the lower type's indifference curve. This cost can be interpreted as the marginal cost of inducing additional productive distortions into the economy, in order to reduce the mimicking rents that are obtainable by higher-type agents. We label it $DC\,(c, y, \theta)$:

$$DC\,(c, y; \theta) \equiv \frac{\tau\,(c, y; \theta)}{u_c\left(c, y; \theta'\right)(1 - \tau\,(c, y; \theta)) + u_y\left(c, y; \theta'\right)} \tag{31}$$

where $\theta' = \min\left\{\widetilde{\theta} \in \Theta : \widetilde{\theta} > \theta\right\}$.[27][28]

The object in the numerator here is the effective marginal tax rate levied on the agent of type θ, which is what the policymaker foregoes for every unit by which the production of that agent is reduced. The object in the denominator is the number of units by which the utility of mimicking types is changed for every unit increase in production for those reporting θ, given that the perturbation takes place along the indifference curve of type θ, whose slope is equal to $(1 - \tau)$. Its inverse is thus the number of units by which production must be reduced in order to reduce minicking utility by one unit.

Note that the single crossing condition implies the denominator is always positive here. Thus higher values for $DC\,(c, y; \theta)$ correspond to higher effective marginal tax rates on agents of type θ. For this reason, the best interpretation

[27] If θ is the maximal element in Θ we can arbitrarily define $u_c\,(c, y; \theta') = 1$ and $u_y\,(c, y; \theta') = 0$ irrespective of c and y. This is for completeness only: we already know that there will be no distortion at the top, in the sense that $\tau\,(c, y; \theta) = 0$, in this case. It therefore makes sense to fix $DC\,(c, y; \theta)$ to zero too.

[28] We will sometimes use the shorthand $u\left(\widehat{\theta}^n; \theta^{n+1}\right)$ to denote the utility of an agent whose true type is θ^{n+1} but who reports a type one lower.

of $DC(c, y; \theta)$ is as an efficiency cost of distorting outcomes. The larger it is, the greater are the deviations from full productive efficiency that the policymaker is willing to tolerate – presumably because these distortions are necessary in order to engineer greater equality across types. Indeed, based on this interpretation we are able to show below that there is a remarkably direct 'efficiency-equity trade-off' at the heart of this class of models.

A full consideration of the marginal resource costs associated with general within-period changes in utilities yields the following, the proof of which is in the appendix:

Proposition 6 ***Intratemporal optimality (iid case):*** *Suppose type draws are iid across agents and time. Then for all time periods $t \geq 1$ and all reporting histories $\widehat{\theta}^t$, the optimal allocation matrix $X_{t+1}^*\left(\widehat{\theta}^t\right)$ satisfies the following condition:*

$$\sum_{\theta^n \in \Theta \backslash \theta^N} \pi_\Theta\left(\theta_t^n\right)\left(\nu_{n+1}-\nu_n\right) DC\left(\theta_t^n\right)=\sum_{\theta^n \in \Theta} \pi_\Theta\left(\theta_t^n\right) \nu_n MC\left(\theta_t^n\right) \tag{32}$$

where ν_n is the nth element of any vector ν that satisfies:

$$\sum_{n=1}^N \pi_\Theta\left(\theta^n\right) \nu_n=0$$

The most useful insights from this condition come when one makes specific choices for the ν vector. Perhaps the most obvious is to limit the $(\nu_{n+1}-\nu_n)$ term to be zero for all n except one. To this end, suppose we pick some $m \in \{1, ..., N-1\}$ and let $\nu_n = -1$ for all $n \leq m$ and $\nu_n = \left[\pi_\Theta\left(\theta > \theta^m\right)\right]^{-1} - 1$ for all $n > m$. By construction the *ex-ante* expected value of ν_n is zero, and we can state the following corollary:

Corollary 2 *Suppose type draws are iid across agents and time. Then for all time periods $t \geq 1$ and all reporting histories $\widehat{\theta}^t$, the optimal allocation matrix $X_{t+1}^*\left(\widehat{\theta}^t\right)$ satisfies the following condition:*

$$\frac{\pi_\Theta\left(\theta_t^m\right)}{\pi_\Theta\left(\theta_t>\theta_t^m\right)} DC\left(\theta_t^m\right)=E\left[MC\left(\theta_t\right) | \theta_t>\theta_t^m\right]-E\left[MC\left(\theta_t\right)\right] \tag{33}$$

E here is the standard expectations operator, being taken under the unique within-period distribution across types, given by π_Θ. Expressed in this way, we have a relationship between the expected marginal cost of utility provision across *all* types at t (the second term on the right-hand side), and the expected marginal cost conditional upon type being higher than θ_t^m (the first term on the right-hand side). The condition states that the higher is the gap between these two expected marginal costs, the more willing the policymaker should be to distort the productive activity of an agent of type θ_t^m. Recall that in general a utilitarian policymaker would like to redistribute utility to those for whom the cost of providing it is lower. Indeed, at an optimal allocation the marginal cost of utility provision would be equalised across all agents. When types are unobservable the cost of providing utility is instead lower for those whose types are relatively low, as a by-product of the need to incentivise higher-type agents. If one were to reduce the utility of higher-type agents any further, this could be done only by distorting (further) the productive activities of those whose types are lower – engineering a movement along their within-period indifference curve, in terms of the graphical analysis, and thus reducing the rents from mimicking. The higher the relative marginal cost of utility provision for high types, the greater the productive distortion the policymaker will be willing to tolerate. This is what (33) is capturing.

Recall that $DC(\theta_t^m)$ is given by:

$$DC(\theta_t^m) \equiv \frac{\tau(\theta_t^m)}{u_c\left(\widehat{\theta}_t^m;\theta_t^{m+1}\right)(1-\tau(\theta_t^m)) + u_y\left(\widehat{\theta}_t^m;\theta_t^{m+1}\right)} \tag{34}$$

So when $DC(\theta_t^m)$ is higher, the effective labour income tax rate on the mth individual is higher too. This highlights an important intuitive point, often missed in popular discussions of optimal tax rates: higher *marginal* rates on lower-type agents are the means by which the incomes of those higher up in the distribution are pulled down. If the right-hand side of (33) is higher, a higher value of $\tau(\theta_t^m)$ can be justified, precisely because it will do more to reduce the gap between the conditional and unconditional expectations of $MC(\theta_t)$. What matters for an individual's welfare is not the marginal rate at which they are able to *transform* leisure into consumption, but the total quantities of these goods that they are able to enjoy. If higher marginal rates on agents of type θ_t^m simply serve to raise additional resources from those higher up in the distribution, and these resources are then used to raise the living standards of all types, then they are evidently in the best interests of θ_t^m types. The greater the measure of individuals whose type is higher than θ_t^m, the greater will be the benefits from increasing the tax rate on θ_t^m in order to reduce the rents of those higher up. This is captured by the $\pi_\Theta(\theta_t > \theta_t^m)$ term in the denominator of the fration on the left-hand side.

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.[29] Though these static models tend to assume a continuous type distribution, we show in a companion paper that the limit of (33) 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

[29] Since (33) is derived by considering perturbations that preserve incentive compatibility within a single time period, it is equally necessary for optimality when the number of time periods is limited to just one.

the arguments in the current paper are equally useful in clarifying the analysis and implementation of static optimal income tax problems as of dynamic ones.

5.5.3 Example with log consumption

A particularly clear version of condition (33) arises when utility is separable, and depends on the log of consumption:

$$u(c, y; \theta) = \ln(c) - \frac{\kappa}{\psi}\left(\frac{y}{e^{\theta}}\right)^{\psi} \tag{35}$$

κ and ψ are parameters, with $\psi - 1$ the inverse of the Frisch elasticity of labour supply, and we let θ implicitly be the log of the marginal product of labour to aid later arguments.Then $DC(\theta_t^m)$ is given by:

$$DC(\theta_t^m) = \frac{\tau(\theta_t^m)}{\kappa y(\theta_t^m)^{\psi-1}\left[e^{-\psi\theta_t^m} - e^{-\psi\theta_t^{m+1}}\right]} \tag{36}$$

and $MC(\theta_t^m)$ is simply equal to $c(\theta_t^m)$. Although we have established the arguments here under an assumption that Θ is a discrete set, the most compact statements result when we let types become arbitrarily close to one another – justified on the grounds that a companion paper confirms the validity of the method when Θ is an interval of $\mathbb{R}$. We then have:

$$\begin{aligned} \lim_{\theta_t^{m+1}\to\theta_t^m}\left[\pi_\Theta(\theta_t^m) DC(\theta_t^m)\right] &= p_\Theta(\theta_t^m)\frac{\tau(\theta_t^m)}{\psi\kappa y(\theta_t^m)^{\psi-1}e^{-\psi\theta_t^m}} \\ &= p_\Theta(\theta_t^m)\frac{\tau(\theta_t^m)}{1-\tau(\theta_t^m)}\frac{c(\theta_t^m)}{\psi} \end{aligned} \tag{37}$$

where we define $p_\Theta(\theta_t^m)$ as the standard density function associated with π_Θ in the continuous-type case. (33) is then:[30]

$$\frac{\tau(\theta_t^m)}{1-\tau(\theta_t^m)}\frac{c(\theta_t^m)}{\psi}\frac{p_\Theta(\theta_t^m)}{\pi_\Theta(\theta_t > \theta_t^m)} = E\left[c(\theta_t)\,|\theta_t > \theta_t^m\right] - E\left[c(\theta_t)\right] \tag{38}$$

So tax rates should be higher the greater is the difference between the expected level of consumption above θ_t^m and its population mean, but lower the higher is the elasticity of labour supply, the hazard rate ($\frac{p_\Theta(\theta_t^m)}{\pi_\Theta(\theta_t > \theta_t^m)}$), or the consumption level of the agent of type θ_t^m. Again, the role of marginal taxes in mitigating the costs that the policymaker attaches to inequality are clearly seen here. The dependence of optimal tax rates on the hazard rate and labour supply elasticity are again both well known from the static case,[31] but have not previously been collected in so compact a statement.

[30] We have used the definition of the labour tax wedge here:

$$\kappa\left(\frac{y(\theta_t^m)}{e^{\theta_t^m}}\right)^{\psi-1}\frac{1}{e^{\theta_t^m}} = (1-\tau(\theta_t^m))\frac{1}{c(\theta_t^m)}$$

[31] Saez (2001) places particular emphasis on the role of the hazard rate in influencing optimal rates at the top of the income distribution.

5.5.4 Discussion: A complete characterisation

There will, in general, exist $N-1$ linearly independent ν vectors satisfying the requirements of Proposition 6. Corollary 2 gives the $N-1$ possibilities for which $\nu_n = -1$ for $n \leq m$, and $\nu_n = [\pi_\Theta(\theta_t > \theta_t^m)]^{-1} - 1$ for $n > m$, for each $m \in \{1, ..., N-1\}$ and any $t \geq 1$. When we move to the non-iid case the set of admissible ν vectors is reduced in an important way, and the simpler optimality condition stated in the Corollary will no longer go through. But for ν vectors that *do* remain admissible, (32) will continue to apply even in that more general setting. This is the chief advantage in stating it in its general form.

Note that in addition to these $N-1$ (linearly distinct) conditions there will always exist an Nth 'intratemporal' requirement: the 'no distortion at the top' result. Further, for all time periods after the first, the generalised inverse Euler condition, (26), 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: one for each adjacent pair of types. 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). Together these restrictions ensure 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.[32]

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 object $MC(\theta_{t-1})$ must, therefore, be sufficient to summarise the past completely.

5.5.5 An 'efficiency-equity trade-off'

There is one final simple and insightful expression can be obtained from Proposition 6 alone. To illustrate it, and without loss of generality, it helps to suppose

[32] For T periods there will generally be $\sum_{t=1}^{T} 2N^t$ variables to determine: a consumption and output allocation for agents with each history at each point in time.

that the N elements of $\Theta \subset \mathbb{R}$ are evenly spaced, so that $\theta^{n+1} - \theta^n = \varepsilon$ for all $n \in \{1, ..., N-1\}$ and some $\varepsilon > 0$.[33] Then suppose that we wish to implement some utility perturbation vector ν whose nth element takes the form:

$$\nu_n = \frac{1}{\varepsilon} \left\{ \theta^n - \sum_{\theta^m \in \Theta} \pi_\Theta (\theta^m) \theta^m \right\} \tag{39}$$

This clearly satisfies $\sum_{\theta^n \in \Theta} \pi_\Theta (\theta^n) \nu_n = 0$, so it provides an admissible vector by which we can augment utilities within a time period at the margin, under the maintained iid assumption.

For the specified ν vector, condition (32) becomes:[34]

$$\varepsilon \cdot E [DC (\theta_t)] = Cov (\theta_t, MC (\theta_t)) \tag{40}$$

where the expectations and covariance operators are taken under the unique distribution π_Θ, across types with a common shock history prior to t.

This is a remarkably clear statement of the 'efficiency-equity trade-off' that is so often at the heart of optimal choice in public economics. Neglecting the normalisation factor ε,[35] the term on the left-hand side gives the average value of the marginal resource cost the policymaker is willing to endure in order to hold down the utility rents that are enjoyed by high-type agents. This cost manifests itself, as we have seen, as a within-period restriction on the productive efficiency of lower-type agents, via a strictly positive within-period effective income tax wedge. Higher values for $DC (\theta_t^n)$ imply higher productive distortions. The term on the right-hand side tracks the degree of cross-sectional inequality in the economy. The marginal cost of providing utility to type θ_t^n in an incentive-compatible way is increasing in type: it is more costly to improve the lot of those who are already doing well. The associated covariance between θ_t^n and $MC (\theta_t^n)$ will be greater the greater is the degree of inequality in welfare across θ_t^n draws. So the overall expression (40) simply states that policymakers should be willing to tolerate inequality only to the extent that the resource costs of doing so are not too great.

---

[33] To see that this is without loss of generality, note that we could always define $\widehat{\Theta}$ as the integers $\{1, ..., N\}$, and the mapping $\theta : \widehat{\Theta} \to \Theta$ by $\theta (n) = \theta^n$, where θ^n is the nth entry in Θ as usual. Then all of the foregoing analysis could be carried out for the utility function $\widehat{u} : \mathbb{R}_+ \times \mathbb{R}_+ \times \widehat{\Theta} \to \mathbb{R}$ defined by $\widehat{u} (c, y; n) = u (c, y; \theta^n)$.

[34] We use the 'no distortion at the top' result here, implying that $DC \left(\theta_t^N\right) = 0$.

[35] This features because $DC (\theta^n)$ is the marginal cost of reducing by one unit the within-period utility an agent of type θ^{n+1} can obtain by mimicking θ^n, holding constant the utility of truth-tellers. In general the further apart θ^n and θ^{n+1} become the lower this cost becomes: as the types become more distinct from one another it is easier to engineer distinct utility changes for them. ε effectively corrects for this effect.

Note that since $u_c (\theta_t^n) (1 - \tau (\theta_t^n)) + u_y (\theta_t^n) = 0$ we have:

$$\varepsilon \cdot DC (\theta_t^n) = \frac{\tau (\theta_t^n)}{\frac{u_c\left(\widehat{\theta}_t^n; \theta_t^{n+1}\right) - u_c (\theta_t^n)}{\varepsilon} (1 - \tau (\theta_t^n)) + \frac{u_y\left(\widehat{\theta}_t^n; \theta_t^{n+1}\right) - u_y (\theta_t^n)}{\varepsilon}}$$

So the left-hand side of (40) will remain well-defined as types become arbitrarily close to one another, the objects in the denominator here approaching the cross-partials $u_{c\theta}$ and $u_{y\theta}$.

To make matters still more concrete, suppose we were again to assume the simple specification for the utility function given by (35). Letting $\varepsilon \to 0$ to approximate the continuous-type case,[36] (40) then becomes:

$$E\left[\frac{\tau\left(\theta_t^n\right)}{1-\tau\left(\theta_t^n\right)}\frac{c\left(\theta_t^n\right)}{\psi}\right] = Cov\left(\theta_t^n, c\left(\theta_t^n\right)\right) \tag{41}$$

The term on the right-hand side in (41) is now a more or less direct measure of (consumption) inequality across types with a shared history, whilst the term on the left-hand side is a simple interaction between the tax rate, consumption level and Frisch elasticity of labour supply.[37] Unsurprisingly, higher taxes in general should be tolerated the greater is inequality, and for any given level of inequality taxes should be higher the less elastic is labour supply (i.e., the higher is ψ). The consumption term on the left-hand side captures the fact that those who are more prosperous on the whole find working relatively easy, and so the resources lost by deterring their production through higher marginal taxes are relatively substantial.

Moreover, since condition (40) would have to hold at the optimum for any *static* income tax problem, it provides an obvious simple restriction for testing the extent to which different tax systems in existence around the world are over-emphasising efficiency or equity considerations relative to the first-best (utilitarian) optimum. This demonstrates the basic point that even if one does not accept the *normative* case for the complex dynamic dependencies allowed in the full model of this paper, there remain *positive* payoffs to our perturbation approach.

5.5.6 Optimal effective income tax rates

The results of the earlier analysis also allow us to demonstrate a further quite general result with important economic implications, requiring no assumptions at all on the evolution of productivity through time. The proof is in the appendix.

Proposition 7 ***Non-negative income taxes:*** *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 $\tau\left(\theta_t^n\right)$ satisfies $\tau\left(\theta_t^n\right) \geq 0$.*

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 Proposition. 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

[36] Again, the discussion is necessarily heuristic here, since the foregoing arguments were not established when Θ is a continuum. Brendon (2011) shows that the optimality conditions explored here do indeed extend to that case in their limiting form.

[37] The Frisch elasticity is $(\psi - 1)^{-1}$.

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 2 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. If the analysis is to aspire to practical implementability 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 easily 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 $\Delta : \mathbb{R} \to \mathbb{R}^{2N}$ and $\Delta_{-1} : \mathbb{R} \to \mathbb{R}^2$ 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 becomes harder to satisfy when probability distributions over future states become type-specific. 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). We can state the following:

Proposition 8 ***No extra distortions in first period and at the top:*** *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 optimal allocation matrix $X_{t+1}^*\left(\theta^t\right)$ satisfies the following condition:*

$$\sum_{\theta^n \in \Theta \backslash \theta^N} \pi_\Theta\left(\theta_{t+1}^n | \theta_t^N\right)\left(\nu_{n+1}-\nu_n\right) DC\left(\theta_{t+1}^n\right)=\sum_{\theta^n \in \Theta} \pi_\Theta\left(\theta_{t+1}^n | \theta_t^N\right) \nu_n MC\left(\theta_{t+1}^n\right) \tag{42}$$

where ν_n is the nth element of any vector ν that satisfies the equation:

$$\sum_{\theta^n \in \Theta} \pi_\Theta\left(\theta_{t+1}^n | \theta_t^N\right) \nu_n=0$$

Similarly in period 1 the optimal allocation matrix X_1^* satisfies the following condition:

$$\sum_{\theta^n \in \Theta \backslash \theta^N} \pi_\Theta\left(\theta_1^n\right)\left(\nu_{n+1}-\nu_n\right) DC\left(\theta_1^n\right)=\sum_{\theta^n \in \Theta} \pi_\Theta\left(\theta_1^n\right) \nu_n MC\left(\theta_1^n\right) \tag{43}$$

where ν_n is the nth element of any vector ν that satisfies the equation:

$$\sum_{\theta^n \in \Theta} \pi_\Theta\left(\theta_1^n\right) \nu_n=0$$

The proof of these claims merely repeats the logic contained in Proposition 6, 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 relatively simple perturbations – 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 more general perturbation arguments 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 $\pi_\Theta\left(\cdot|\theta_t^n\right)$ and the 'mimicker's' measure $\pi_\Theta\left(\cdot|\theta_t^{n+1}\right)$. In the iid case we were able implement any vector of marginal utility increments across agents at $t+1$ provided this vector satisfied $\sum_{\theta^n\in\Theta}\pi_\Theta\left(\theta^n\right)\nu_n = 0$ for the unique probability measure across $t+1$ types, π_Θ. In general one can always find $N-1$ linearly independent ν vectors that satisfy this condition.

When shocks are Markov the probability measure across $t+1$ types is no longer unique. It depends on the agent's true type at t. But we can preserve incentive compatibility for both truth-tellers and *relevant* (i.e., 'one-higher') mimickers provided we perturb utilities at the margin according to a vector ν that *jointly* satisfies the *two* conditions:

$$\sum_{\theta^m\in\Theta}\pi_\Theta\left(\theta_{t+1}^m|\theta_t^n\right)\nu_m = \sum_{\theta^m\in\Theta}\pi_\Theta\left(\theta_{t+1}^m|\theta_t^{n+1}\right)\nu_m = 0 \qquad (44)$$

In general one can always find $N-2$ linearly independent ν vectors for which this condition is satisfied. For this reason the movement to Markov probabilities will ultimately have denied us the capacity to carry out intratemporal perturbations in precisely one dimension.

Lemma 1 can now be easily adjusted to cover intratemporal perturbations in the Markov case:

Lemma 2 *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 (44) it is possible to perturb the optimal allocations $X_{t+1}^*\left(\theta^t\right)$ 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. 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.

The required intratemporal optimality conditions can now be stated formally:

Proposition 9 ***Intratemporal optimality (Markov case):*** *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 optimal allocation matrix $X_{t+1}^*\left(\theta^t\right)$ satisfies the following condition:*

$$\sum_{\theta^m \in \Theta \backslash \theta^N} \pi_\Theta\left(\theta_{t+1}^m | \theta_t^n\right)\left(\nu_{m+1} - \nu_m\right) DC\left(\theta_{t+1}^m\right) = \sum_{\theta^m \in \Theta} \pi_\Theta\left(\theta_{t+1}^m | \theta_t^n\right) \nu_m MC\left(\theta_{t+1}^m\right) \tag{45}$$

where ν_m is the mth element of any vector ν that satisfies the two restrictions:

$$\sum_{\theta^m \in \Theta} \pi_\Theta\left(\theta_{t+1}^m | \theta_t^n\right) \nu_m = \sum_{\theta^m \in \Theta} \pi_\Theta\left(\theta_{t+1}^m | \theta_t^{n+1}\right) \nu_m = 0$$

The proof again follows directly from earlier arguments so is omitted here. Together with the '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),[38] and there are $N-1$ binding incentive compatibility constraints. From a purely analytical perspective 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 of the problem will be to provide this missing equation, and this is done in the next sub-section.

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

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 2. 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' 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 2 to confirm incentive compatibility for dynamic perturbations. We have the following, the proof of which is in the appendix:

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 the two restrictions:*

$$\sum_{\theta^m \in \Theta} \pi_\Theta\left(\theta_{t+1}^m|\theta_t^n\right)\nu_m = 0$$

and

$$\sum_{\theta^m \in \Theta} \pi_\Theta\left(\theta_{t+1}^m|\theta_t^{n+1}\right)\nu_m = 1$$

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 compati-

bility 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 $\nu_m \delta$ 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. The proof is in the appendix.

Proposition 10 ***Dynamic cost-spreading:*** *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 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:*

$$\begin{aligned} & \beta R_{t+1} DC\left(\theta_t^n\right) \\ = & \sum_{\theta^m \in \Theta \backslash \theta^N} \pi_\Theta\left(\theta_{t+1}^m | \theta_t^n\right)\left(\nu_{m+1} - \nu_m\right) DC\left(\theta_{t+1}^m\right) - \sum_{\theta^m \in \Theta} \pi_\Theta\left(\theta_{t+1}^m | \theta_t^n\right) \nu_m MC\left(\theta_{t+1}^m\right) \end{aligned} \tag{46}$$

where ν_m is the mth element of any vector ν that satisfies the two restrictions given in Lemma 3.

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 (2011) in the context of dynamic Mirrleesian problems. Equation (46) 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.[39] 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. As Proposition 8 shows, if $t = 1$ the right-hand side of (46) would be zero, so in all subsequent periods an 'uncommitted' policymaker 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 seems likely to result in greater equality at $t+1$ the higher is the marginal tax rate for an agent at t. This is because, as discussed, higher marginal rates are really a means

[39] 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).

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, 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. The next subsection confirms that this is indeed the case, and in the process provides a novel expression for the dynamics of the labour wedge, and of the policymaker's willingness to trade off equity and efficiency through time.

6.3 Equity, efficiency, and the dynamics of the labour wedge

Again, it seems desirable to arrive at some more direct economic results as payoff for the abstract analysis above. The most useful statements can be obtained by generalising to the Markov setting our earlier treatment of the equity-efficiency trade-off. To this end, suppose once more (and again, without loss of generality) that the elements of $\Theta \subset \mathbb{R}$ are evenly spaced on the real line, with $\theta^{n+1} - \theta^n = \varepsilon$ for all $n \in \{1, ..., N-1\}$. Then for any $t \geq 1$ and any agent whose productivity in period t was θ_t^m for $m \in \{1, ..., N-1\}$ we can fix a utility increment vector ν for application at $t+1$ whose nth entry is given by:

$$\nu_n = \frac{\left\{\theta_{t+1}^n - \sum_{\theta_{t+1} \in \Theta} \pi_\Theta \left(\theta_{t+1} | \theta_t^m\right) \theta_{t+1}\right\}}{\sum_{\theta_{t+1} \in \Theta} \pi_\Theta \left(\theta_{t+1} | \theta_t^{m+1}\right) \theta_{t+1} - \sum_{\theta_{t+1} \in \Theta} \pi_\Theta \left(\theta_{t+1} | \theta_t^m\right) \theta_{t+1}} \tag{47}$$

The denominator here is the difference in expected type in period t between those who drew θ_t^{m+1} and those who drew θ_t^m in the previous period, whilst the numerator is the difference between θ_{t+1}^n and its expected value for the previous period's truth-tellers. It is clear by inspection that this vector will satisfy the requirements of a ν vector in Proposition 10; re-writing condition (46) then gives:

$$\beta R_{t+1} \varepsilon DC\left(\theta_t^m\right) = \frac{\varepsilon \cdot E\left[DC\left(\theta_{t+1}\right) | \theta_t^m\right] - Cov\left[\left(\theta_{t+1}, MC\left(\theta_{t+1}\right)\right) | \theta_t^m\right]}{\frac{1}{\varepsilon}\left\{E\left[\theta_{t+1} | \theta_t^{m+1}\right] - E\left[\theta_{t+1} | \theta_t^m\right]\right\}} \tag{48}$$

where the conditional expectation and covariance terms are taken across types with common productivity draws to $t-1$.[40]

[40] Farhi and Werning (2011) provide a special case of this condition under the assumptions that consumption and labour supply are separable in preferences, the disutility of work is isoelastic, and productivity is AR(1).

ε here should again be treated as a normalising factor, allowing for the fact that the cost of incentivising types 'one-higher' than truth-tellers through movements along a truth-teller's indifference curve (i.e., $DC(\theta)$) will become harder the smaller the difference between truth-tellers and mimickers.[41] The term on the left-hand side of the equation is then proportional to the marginal efficiency cost of the labour tax wedge imposed in period t on type θ_t^m. Proposition 7 implies this object will always be non-negative. The term on the right-hand side can be digested in pieces. The fraction's numerator is familiar from equation (40) above: it gives the difference between the average marginal efficiency costs of the within-period tax distortions implemented in period $t+1$, and the costs from excessive inequality. The denominator, meanwhile, is the difference in the expected type draw for $t+1$ between agents whose t-period draws are θ_t^{m+1} and θ_t^m respectively (relative to ε, the absolute difference between these types). It is a natural measure of the persistence of productivity draws, and one would expect it to be positive in all cases of interest.

Taken together, (48) has a remarkably clear message: persistent shocks mean the policymaker should implement more equality for any given within-period efficiency cost. This is what it means for the numerator of the fraction to be positive, which – according to (48) – it must be under any optimal scheme whenever the tax wedge was strictly positive in the prior period. Note that this relatively high equality only holds across agents with a common shock history up to period t; it need not carry over to the economy as a whole. Nonetheless, it is behind the numerical result obtained under separable preferences by Farhi and Werning (2011) that optimal tax wedges should drift upwards on average over time when type draws follow a random walk.

Interestingly Farhi and Werning focus on the cross-sectional profile of the labour wedge to argue that the extra dynamic complications introduced by (48) result in more *regressive* outcomes relative to the iid case – since marginal taxes tend to drift upwards most over time for those who receive relatively low type draws. The interpretation we provide here appears quite the opposite: persistent shocks introduce a bias towards equity relative to efficiency, which is hardly a 'regressive' move. Reconciling the two interpretations highlights an important point often overlooked in popular discussions of tax reform: high *marginal* tax rates at any given point in the income distribution are a way for *average* rates to be increased higher up. The whole reason for inducing within-period labour supply distortions in the first place is to reduce the compensation that more productive agents must be paid to incentivise them to work. Hence the regressivity in marginal rates identified by Farhi and Werning is really just a means for engineering greater progressivity in utility outcomes in equilibrium.

Two other interesting observations can be made about (48). First, note that it has a 'reset' feature. If at any point in time t the highest productivity θ_t^N is drawn then we know from Proposition 8 that the set of optimality conditions at $t+1$ becomes identical to the iid case. This implies in particular that equation (40) will hold: the numerator of the fraction on the right-hand side of (48) will

[41] See the discussion in Section 5.5.5.

equal zero, with efficiency and equity considerations exactly balancing. Any average drift towards greater equality always has the potential to be dominated by 'no distortion at the top'.

Second, if type draws are I(1) and the interest rate is equal to the inverse of the discount factor then the object $DC(\theta)$ will in general increase over time. In this event the denominator of the fraction on the right-hand side is one, and since the covariance term will be (at least weakly) positive the equation reduces to a random walk plus a stochastic drift term. Thus full persistence in type draws translates into persistence *plus drift* in labour wedges. This is a generalisation of the 'tax smoothing' result of Golosov, Tsyvinski and Werning (2006), who showed that labour wedges should be constant through time in the event that type draws are drawn 'once and for all' in the first period of the model. In that case there is no variation in type draws at $t+1$ for any given draw at t, implying the covariance term drops from (48) whilst the conditional expectations reduce to certainties. Allowing continued uncertainty instead biases us away from tax smoothing and towards continued upwards drift in the labour tax wedge.

As a point of comparison with the literature it is again instructive briefly to note a special case of (48) in which $\varepsilon \to 0$ and more specific preference assumptions are made. To this end we once more let the utility function take the form:

$$u(c, y; \theta) = \ln(c) - \frac{\kappa}{\psi}\left(\frac{y}{e^{\theta}}\right)^{\psi} \tag{49}$$

The optimality condition is then:

$$\beta R_{t+1} \frac{\tau(\theta_t)}{1-\tau(\theta_t)} \frac{c(\theta_t)}{\psi} = \frac{E\left[\frac{\tau(\theta_{t+1})}{1-\tau(\theta_{t+1})} \frac{c(\theta_{t+1})}{\psi} |\theta_t\right] - Cov\left[(\theta_{t+1}, c(\theta_{t+1})) |\theta_t\right]}{\frac{\partial}{\partial \theta_t} E\left[\theta_{t+1}|\theta_t\right]} \tag{50}$$

This expression is, reassuringly, identical to one already provided by Farhi and Werning (2011) under the same preference restrictions;[42] but we have arrived at it by a very different route. The general expression from which these authors take (50) as a special case is a comparatively unwieldy object relative to (48), dependent in particular on the derivative of a weighting function that is defined in their paper, and on a costate variable that their method does not allow them to eliminate when preferences are non-separable. The method presented in this paper appears to permit much sharper intuition regarding the general character of optimal distortions through time.

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

[42] C.f. their equation (13).

equal to the inverse of the discount factor β. Then the generalised inverse Euler equation can be written as:

$$\frac{1-\alpha\left(\theta_t\right)}{u_c\left(\theta_t\right)+u_y\left(\theta_t\right)\alpha\left(\theta_t\right)}=\sum_{\theta_{t+1}\in\Theta}\pi_\Theta\left(\theta_{t+1}|\theta_t\right)\frac{1-\alpha\left(\theta_{t+1}\right)}{u_c\left(\theta_{t+1}\right)+u_y\left(\theta_{t+1}\right)\alpha\left(\theta_{t+1}\right)} \tag{51}$$

That is to say, we have a martingale in the marginal cost of (locally incentive compatible) utility provision, which we have chosen to write out in full here. 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 policymaker'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 (51) – the marginal cost of utility provision – for preference structures more general than the separable case. A first step is the following.

Lemma 4 *Under an optimal plan that solves the restricted problem,* $u_c\left(\theta_t\right)+u_y\left(\theta_t\right)\alpha\left(\theta_t\right)>0$ *always holds.*

Proof. By definition

$$\alpha\left(\theta_t^n\right)=\frac{u_c\left(\theta_t^n\right)-u_c\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)}{u_y\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)-u_y\left(\theta_t^n\right)} \tag{52}$$

for $n<N$, and $\alpha\left(\theta_t^N\right)=0$. In the latter case the result follows immediately from $u_c\left(\theta_t\right)>0$. In the former case we have from equation (5):

$$\frac{u_y\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)}{u_y\left(\theta_t^n\right)}<\frac{u_c\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)}{u_c\left(\theta_t^n\right)} \tag{53}$$

[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.

Rewriting our object of interest, we have:

$$u_c(\theta_t) + u_y(\theta_t)\,\alpha(\theta_t) = \frac{u_c\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right) - u_c(\theta_t^n)\,\frac{u_y\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)}{u_y(\theta_t^n)}}{1 - \frac{u_y\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)}{u_y(\theta_t^n)}} \tag{54}$$

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)). ■

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 5 $\frac{1-\alpha(\theta_t)}{u_c(\theta_t)+u_y(\theta_t)\alpha(\theta_t)} > 0$ *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 4. 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 11 ***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):

$$u(c, y; \theta) = \frac{c^{1-\varsigma}}{1-\varsigma} \exp\left\{(\varsigma - 1) v\left(\frac{y}{\theta}\right)\right\} \tag{55}$$

with the labour disutility schedule v defined by:

$$v(l) = \frac{l^{1+v}}{1+v} \tag{56}$$

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

[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.

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:

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

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 the relatively high value: $\varsigma = 10$. For the other parameters we choose values $\upsilon = 2$ 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 4 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

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

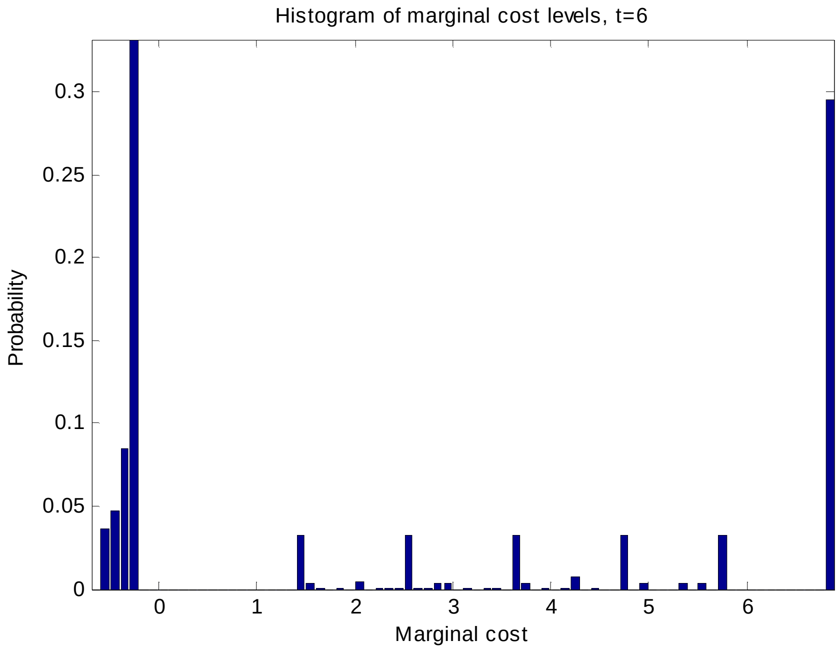

Figure 4: Distribution of marginal cost of utility provision in 6th period

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.

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 reducing 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 5 – 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 12 ***Deterred savings (3):*** *Suppose type draws are iid across agents and time. 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 (27), 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:

$$\frac{1-\alpha\left(\theta_t\right)}{u_c\left(\theta_t\right)+u_y\left(\theta_t\right)\alpha\left(\theta_t\right)} > 0 \tag{57}$$

holds – which we now know to be the case under complementarity and iid productivity draws by Lemma 5. 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 (27). 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.

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 a functional operator.

But the focus of the paper has been on exploiting the analytical results that a perturbation approach can expose, and here there are several. First, we have shown that many of the well-known results from static income tax theory generalise to the dynamic case. In particular, 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, and agents whose type is the highest always have a zero effective marginal tax rate. We have also been able to obtain a set of optimality conditions that is considerably simpler and easier to interpret than those conventionally analysed in the static income tax literature, linking effective marginal income tax rates at each point in the distribution to a small number of economically meaningful variables. Thus it is hoped that the interaction between this paper and the static optimal tax literature could work both ways.

In general higher effective income tax rates correspond to higher productive distortions, and are undesirable from a policymaker's perspective for this reason. But against this 'efficiency' cost must be balanced the 'equity' benefit that higher tax rates deliver, reducing the rents that more productive agents are able to enjoy. Optimal policy can be interpreted as resolving a trade-off between these considerations, given the utilitarian objective that is assumed. In this paper we are able to provide a novel summary statement of this trade-off, showing how the expected value of the resource cost from productive distortions across types with a common prior history should be linked to a measure of the covariance between agents' types and the marginal cost of providing utility to them. Since this marginal cost will be higher the greater the quantity of utility an agent already enjoys, this covariance term is interpretable as a measure of inequality. We show that the higher it is, the higher the average productive distortion should be.

Perhaps more surprisingly, we also show that when productivities are persistent through time the policymaker should be willing to tolerate more and more productive inefficiency (across types with a common history), relative to the degree of inequality. The reason for this derives from a desire to reduce the returns from mimicking. The associated optimality condition suggests that provided persistence is great enough, effective marginal tax rates should drift upwards through time on average. This result has already been obtained numerically, and analytically under specific preference assumtions, by Farhi and Werning (2011). Here we show the extent to which it will generalise.

Turning to optimal 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. This seems an interesting avenue for further research.

9 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] Atkeson, A. and R.E. Lucas Jr. (1992), 'On Efficient Distribution with Private Information', *Review of Economic Studies*, 59, 427–453.

[3] Brendon, C.F. (2011), 'Applying perturbation methods to dynamic tax problems with a continuum of types', Manuscript (available on request).

[4] 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).

[5] da Costa, C., and I. Werning (2002), 'Commodity Taxation and Social Insurance', Manuscript (available at http://economics.mit.edu/files/1265).

[6] Dasgupta, P. (1982), 'Utilitarianism, Information and Rights', in A. Sen and B. Williams (eds.), *Utilitarianism and Beyond*, Cambridge, Cambridge University Press.

[7] Diamond, P.A., and J.A. Mirrlees (1978), 'A Model of Social Insurance with Variable Retirement', *Journal of Public Economics*, 10, 295–336.

[8] 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.

[9] Farhi, E. and I. Werning (2007), 'Inequality and Social Discounting', *Journal of Political Economy*, 115, 365–402.

[10] Farhi, E. and I. Werning (2011), 'Insurance and Taxation over the Life Cycle', Manuscript (available at http://econ-www.mit.edu/files/6588).

[11] Fernandes, A. and C. Phelan (2000), 'A Recursive Formulation for Repeated Agency with History Dependence', *Journal of Economic Theory*, 91, 223–247.

[12] Golosov, M., N. Kocherlakota and A. Tsyvinski (2003), 'Optimal Indirect and Capital Taxation', *Review of Economic Studies*, 70, 569–587.

[13] Golosov, M., M. Troshkin and A. Tsyvinski (2011), 'Optimal Taxation: Merging Micro and Macro Approaches', *Journal of Money, Credit and Banking*, 43, 147–174.

[14] Golosov, M., M. Troshkin and A. Tsyvinski (2011), 'Optimal Dynamic Taxes', Manuscript (available at http://scholar.princeton.edu/golosov/files/odt30.pdf)

[15] 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.

[16] Greenwald, B., and J.E. Stiglitz (1986), 'Externalities in Economies with Imperfect Information and Incomplete Markets', *Quarterly Journal of Economics*, 101, 229–264.

[17] M. Kapička (2011), '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.*

[18] Kocherlakota, N.R. (1996), 'Implications of Efficient Risk Sharing without Commitment', *Review of Economic Studies*, 63, 595-609.

[19] Kocherlakota, N.R. (2005), 'Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation', *Econometrica*, 73, 1587–1621.

[20] Kocherlakota, N.R. (2011), *The New Dynamic Public Finance*, Princeton, Princeton University Press.

[21] Mankiw, N.G., M. Weinzierl and D. Yagan (2009), 'Optimal Taxation in Theory and Practice' *Journal of Economic Perspectives*, 23, 147–174.

[22] Marcet, A. and R. Marimon (1998), 'Recursive Contracts', Economics Working Papers eco98/37, European University Institute.

[23] Marcet, A. and R. Marimon (2011), 'Recursive Contracts', CEP Discussion Paper 1055, Centre for Economic Performance, LSE.

[24] Mele, A. (2011), 'Repeated Moral Hazard and Recursive Lagrangeans', Working Paper (available at http://sites.google.com/site/meleantonio/#research).

[25] Messner, M., N. Pavoni and C.M. Sleet (2011), 'Recursive Methods for Incentive Problems', Manuscript (available at SSRN: http://ssrn.com/abstract=1792816).

[26] Mirrlees, J.A. (1971), 'An Exploration in the Theory of Optimum Income Taxation', *Review of Economic Studies*, 38, 175–208.

[27] 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).

[28] Phelan, C. (2006), 'Opportunity and Social Mobility', *Review of Economic Studies*, 73, 487–505.

[29] 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).

[30] Saez, E. (2001), 'Using Elasticities to Derive Optimal Income Tax Rates', *Review of Economic Studies*, 68, 205–229.

[31] 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).

[32] 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).

[33] 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.

A Appendix

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 $u_{cy}^2 - u_{cc}u_{yy} < 0$,[49] and we have $u_y < 0$, so if $u_{cc} + u_{cy} < 0$ then condition (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 (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

$$W\left(\theta_t^n; \theta_t^n, \widehat{\theta}^{t-1}\right) \geq W\left(\theta_t^m; \theta_t^n, \widehat{\theta}^{t-1}\right)$$

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:

$$\begin{aligned} & W\left(\theta_t^{n-1}; \theta_t^n, \widehat{\theta}^{t-1}\right) - W\left(\theta_t^{n-1}; \theta_t^{n-1}, \widehat{\theta}^{t-1}\right) \\ > \; & W\left(\theta_t^{n-2}; \theta_t^n, \widehat{\theta}^{t-1}\right) - W\left(\theta_t^{n-2}; \theta_t^{n-1}, \widehat{\theta}^{t-1}\right) \end{aligned}$$

[49] Consider movements according to the vector $\begin{bmatrix} -\frac{u_{yc}}{u_{cc}} \\ 1 \end{bmatrix}$ here.

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 3

We consider now a perturbation to outcomes in both time t and time $t+1$. Specifically, we wish to choose $\Delta(\delta)$ and $\Delta_{-1}(\delta)$ 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 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.

Proof. 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:

$$\Delta_{-1}(\delta) = \left(\phi^c\left(-\beta\delta; c_t^*, y_t^*, \theta_t\right), \phi^y\left(-\beta\delta; c_t^*, y_t^*, \theta_t\right)\right) \tag{58}$$

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

$$u\left(c^* + \phi^c(k; c^*, y^*, \theta), y^* + \phi^y(k; c^*, y^*, \theta); \theta\right) = u(c^*, y^*; \theta) + k \tag{59}$$

$$u\left(c^* + \phi^c(k; c^*, y^*, \theta), y^* + \phi^y(k; c^*, y^*, \theta); \theta'\right) = u\left(c^*, y^*; \theta'\right) + k \tag{60}$$

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

$$u\left(c^* + \phi^c(k; c^*, y^*, \theta), y^*; \theta\right) = u(c^*, y^*; \theta) + k \tag{61}$$

$$\phi^y(k; c^*, y^*, \theta) = 0 \tag{62}$$

That is to say, $\phi^c(k; c^*, y^*, \theta)$ and $\phi^y(k; c^*, y^*, \theta)$ 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:

$$\left[\phi^c\left(\delta; c_{t+1}^*, y_{t+1}^*, \theta_{t+1}^n\right), \phi^y\left(\delta; c_{t+1}^*, y_{t+1}^*, \theta_{t+1}^n\right)\right] \tag{63}$$

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:

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

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:

$$\begin{aligned}&\beta\left[\phi_1^c\left(0; c_t^*, y_t^*, \theta_t\right) - \phi_1^y\left(0; c_t^*, y_t^*, \theta_t\right)\right] \\ =\ &R_{t+1}^{-1}\sum_{\theta_{t+1}\in\Theta}\pi_\Theta\left(\theta_{t+1}|\theta^t\right)\left[\phi_1^c\left(0; c_{t+1}^*, y_{t+1}^*, \theta_{t+1}^n\right)\right. \\ &\left.-\phi_1^y\left(0; c_{t+1}^*, y_{t+1}^*, \theta_{t+1}^n\right)\right]\end{aligned} \tag{64}$$

where ϕ_1^c denotes the derivative of ϕ^c with respect to its first argument. By total differentiation of conditions (59) to (62) with respect to k it is easy to show:

$$\phi_1^c\left(0; c^*, y^*, \theta\right) - \phi_1^y\left(0; c^*, y^*, \theta\right) = \frac{1-\alpha\left(c^*, y^*; \theta\right)}{u_c\left(c^*, y^*; \theta\right) + u_y\left(c^*, y^*; \theta\right)\alpha\left(c^*, y^*; \theta\right)} \tag{65}$$

The result follows. ■

A.4 Proof of Proposition 4

The result follows from the arguments already given in the event that condition 1 of the Proposition 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 (3) this implies $u_{c\theta} > 0$. We also know $u_{y\theta} > 0$, so in general we will have $\alpha\left(c, y; \theta\right) \leq 0$, 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:

$$\frac{1-\alpha\left(\theta\right)}{u_c\left(\theta\right) + u_y\left(\theta\right)\alpha\left(\theta\right)} \geq \frac{1}{u_c\left(\theta\right)} \tag{66}$$

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

$$\begin{aligned} & \sum_{\theta_{t+1}\in\Theta} \pi_\Theta \left(\theta_{t+1}|\theta^t\right) \frac{1-\alpha\left(\theta_{t+1}\right)}{u_c\left(\theta_{t+1}\right)+u_y\left(\theta_{t+1}\right)\alpha\left(\theta_{t+1}\right)} \\ \geq & \sum_{\theta_{t+1}\in\Theta} \pi_\Theta \left(\theta_{t+1}|\theta^t\right) \frac{1}{u_c\left(\theta_{t+1}\right)} \\ \geq & \left[\sum_{\theta_{t+1}\in\Theta} \pi_\Theta \left(\theta_{t+1}|\theta^t\right) u_c\left(\theta_{t+1}\right)\right]^{-1} \end{aligned} \tag{67}$$

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 $u_c\left(c_t^*, y_t^*; \theta_t\right) = u_y\left(c_t^*, y_t^*; \theta_t\right)$, so we have:

$$\begin{aligned} R_{t+1}\beta\frac{1}{u_c\left(\theta_t\right)} &= R_{t+1}\beta\frac{1-\alpha\left(\theta_t\right)}{u_c\left(\theta_t\right)+u_y\left(\theta_t\right)\alpha\left(\theta_t\right)} \\ &= \sum_{\theta_{t+1}\in\Theta} \pi_\Theta \left(\theta_{t+1}|\theta^t\right) \frac{1-\alpha\left(\theta_{t+1}\right)}{u_c\left(\theta_{t+1}\right)+u_y\left(\theta_{t+1}\right)\alpha\left(\theta_{t+1}\right)} \\ &\geq \left[\sum_{\theta_{t+1}\in\Theta} \pi_\Theta \left(\theta_{t+1}|\theta^t\right) u_c\left(\theta_{t+1}\right)\right]^{-1} \end{aligned} \tag{68}$$

The result then follows from trivial manipulation.

A.5 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 for all $n \in \{1, ..., N\}$:

$$u\left(c_{n,t}^* + \delta_n^c\left(\delta\right), y_{n,t}^* + \delta_n^y\left(\delta\right); \theta_t^n\right) = u\left(c_{n,t}^*, y_{n,t}^*; \theta_t^n\right) + \nu_n\delta \tag{69}$$

$$u\left(c_{n,t}^* + \delta_n^c\left(\delta\right), y_{n,t}^* + \delta_n^y\left(\delta\right); \theta_t^{n+1}\right) = u\left(c_{n,t}^*, y_{n,t}^*; \theta_t^{n+1}\right) + \nu_{n+1}\delta \tag{70}$$

where $\delta_n^c\left(\delta\right)$ and $\delta_n^y\left(\delta\right)$ are the perturbations to the nth agent's consumption and output levels respectively. For the Nth agent we just need:

$$u\left(c_{N,t}^* + \delta_N^c\left(\delta\right), y_{N,t}^*; \theta_{t+1}^N\right) = u\left(c_{N,t}^*, y_{N,t}^*; \theta_t^N\right) + \nu_N\delta \tag{71}$$

and we normalise $\delta_N^y\left(\delta\right) = 0$.[50]

[50] This is analogous to the normalisation $\phi^y\left(\theta, k; c^*, y^*\right) = 0$ in equation (62).

Equations (69) and (71) 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 (70) 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 – 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. The impact of the perturbations on discounted expected utility from the perspective of prior time periods is left unchanged by the fact that $\sum_{n=1}^{N} \pi_\Theta(\theta^n)\nu_n = 0$, where this probability measure is common to all true types by the iid assumption. This completes the proof.

A.6 Proof of Proposition 6

Our aim is to construct a perturbation schedule $\Delta(\delta)$, to be applied in period t, whose effect effect on the utility of an agent of type θ_t^n will always equal $\nu_n\delta$, and then to consider the marginal impact on the policymaker's resources as δ is moved away from zero, given that this perturbation is applied to the optimal allocation. By Lemma 1 we know such a perturbation can be constructed in a manner that preserves incentive compatibility, through the consumption and output perturbations $\delta_n^c(\delta)$ and $\delta_n^y(\delta)$ that are defined in the proof of that Lemma. The net cost of these perturbations on the policymaker's within-period resources in t (per agent with the relevant prior history) will be given by:

$$\sum_{n=1}^{N} \pi_\Theta(\theta_t^n)\left[\delta_n^c(\delta) - \delta_n^y(\delta)\right]$$

Hence the marginal cost as δ moves away from zero will be:

$$\sum_{n=1}^{N} \pi_\Theta(\theta_t^n)\left[\left.\frac{d\delta_n^c(\delta)}{d\delta}\right|_{\delta=0} - \left.\frac{d\delta_n^y(\delta)}{d\delta}\right|_{\delta=0}\right]$$

The rest of the proof is just a matter of showing that when this object is evaluated at the optimum and set equal to zero, the espression in the Proposition results.

Totally differentiating the restrictions that define $\delta_n^c(0)$ and $\delta_n^y(0)$, one can show:

$$\frac{d\delta_n^c(0)}{d\delta} = \frac{\frac{u_y(\theta_t^n)}{u_c(\theta_t^n)}\frac{\nu_{n+1}}{u_c\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)} - \frac{u_y\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)}{u_c\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)}\frac{\nu_n}{u_c(\theta_t^n)}}{\frac{u_y(\theta_t^n)}{u_c(\theta_t^n)} - \frac{u_y\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)}{u_c\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)}} \tag{72}$$

$$\frac{d\delta_n^y(0)}{d\delta} = \frac{\frac{\nu_n}{u_c(\theta_t^n)} - \frac{\nu_{n+1}}{u_c\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)}}{\frac{u_y(\theta_t^n)}{u_c(\theta_t^n)} - \frac{u_y\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)}{u_c\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)}} \tag{73}$$

for $n \in \{1, ..., N-1\}$, and

$$\frac{d\delta_N^c(0)}{d\delta} = \frac{\nu_N}{u_c\left(\theta_t^N\right)} \tag{74}$$

With some straightforward algebra it is then possible to show for all $n \in \{1, ..., N-1\}$:

$$\begin{aligned} \frac{d\delta_n^c(0)}{d\delta} - \frac{d\delta_n^y(0)}{d\delta} &= \nu_n \frac{1-\alpha(\theta_t^n)}{u_c(\theta_t^n) + \alpha(\theta_t^n) u_y(\theta_t^n)} \\ &\quad - (\nu_{n+1} - \nu_n) \frac{\tau(\theta_t^n)}{u_c\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)(1-\tau(\theta_t^n)) + u_y\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)} \end{aligned} \tag{75}$$

and since we normalise $\delta_N^y(0)$ and $\alpha\left(\theta_t^N\right) = 0$ we additionally have:

$$\begin{aligned} \frac{d\delta_N^c(0)}{d\delta} - \frac{d\delta_N^y(0)}{d\delta} &= \frac{d\delta_N^c(0)}{d\delta} \\ &= \frac{\nu_N}{u_c\left(\theta_t^N\right)} \\ &= \nu_N \frac{1-\alpha(\theta_t^n)}{u_c(\theta_t^n) + \alpha(\theta_t^n) u_y(\theta_t^n)} \end{aligned} \tag{76}$$

The result follows immediately from these statements, together with the definitions of $MC(\theta_t^n)$ and $DC(\theta_t^n)$.

A.7 Proof of Proposition 7

Consider the perturbation given by a movement along the within-period indifference curve of the nth agent, with no changes to the allocations of any other agents. If consumption and output are being jointly reduced this will move us strictly within the constraint set of the relaxed problem, since the net impact on the utility obtainable from reporting the relevant θ^t at t is zero for truth-tellers and strictly negative for 'downwards mimickers' (by single crossing), and expected utility in prior periods is left completely unaffected regardless of the distribution under which it is assessed, by the fact that all agents who remain truth-tellers at t are indifferent to this perturbation. Hence the marginal cost per unit reduction in the utility of potential mimickers must be weakly positive, given that the optimal solution in the relaxed constraint set solves the general problem. From our earlier results, this implies:

$$\frac{\tau(\theta_t^n)}{u_c\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)(1-\tau(\theta_t^n)) + u_y\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)} \geq 0 \tag{77}$$

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

$$\begin{aligned}\left(1-\tau\left(\theta_{t+1}^n\right)\right) &= -\frac{u_y\left(\theta_{t+1}^n\right)}{u_c\left(\theta_{t+1}^n\right)} \\ &> -\frac{u_y\left(\widehat{\theta}_{t+1}^n;\theta_{t+1}^{n+1}\right)}{u_c\left(\widehat{\theta}_{t+1}^n;\theta_{t+1}^{n+1}\right)}\end{aligned} \tag{78}$$

where the last inequality is an application of the single-crossing condition. Hence the denominator in condition (77) will be strictly positive, and the result follows.

A.8 Proof of Lemma 3

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 $m \in \{1,...,N\}$:

$$u\left(c_{m,t+1}^* + \delta_{m,t+1}^c(\delta), y_{m,t+1}^* + \delta_{m,t+1}^y(\delta); \theta_{t+1}^m\right) = u\left(c_{m,t+1}^*, y_{m,t+1}^*; \theta_{t+1}^m\right) + \nu_m\delta \tag{79}$$

$$u\left(c_{m,t+1}^* + \delta_{m,t+1}^c(\delta), y_{m,t+1}^* + \delta_{m,t+1}^y(\delta); \theta_{t+1}^{m+1}\right) = u\left(c_{m,t+1}^*, y_{m,t+1}^*; \theta_{t+1}^{m+1}\right) + \nu_{m+1}\delta \tag{80}$$

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:

$$u\left(c_{N,t+1}^* + \delta_{N,t+1}^c(\delta), y_{N,t+1}^*; \theta_{t+1}^N\right) = u\left(c_{N,t+1}^*, y_{N,t+1}^*; \theta_{t+1}^N\right) + \nu_N\delta \tag{81}$$

and we normalise $\delta_{N,t+1}^y(\delta) = 0$.

The proof of Lemma 1 shows that these conditions can indeed be satisfied by appropriate choice of $\delta_{m,t+1}^c(\delta)$ and $\delta_{m,t+1}^y(\delta)$ 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 $\delta_{n,t}^c(\delta)$ and $\delta_{n,t}^y(\delta)$ schedules that satisfy:

$$u\left(c_{n,t}^* + \delta_{n,t}^c(\delta), y_{n,t}^* + \delta_{n,t}^y(\delta); \theta_t^n\right) = u\left(c_{n,t}^*, y_{n,t}^*; \theta_t^n\right) \tag{82}$$

$$u\left(c_{n,t}^* + \delta_{n,t}^c(\delta), y_{n,t}^* + \delta_{n,t}^y(\delta); \theta_t^{n+1}\right) = u\left(c_{n,t}^*, y_{n,t}^*; \theta_t^{n+1}\right) - \beta\delta \tag{83}$$

Again, by the single crossing condition the utility of the agent of type θ_t^{n+1} 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.9 Proof of Proposition 10

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 ν. By earlier arguments (c.f. proof of Proposition 6), the marginal cost of the $\Delta(\delta)$ perturbation as δ is moved away from 0, assessed from the perspective of time t, will be:

$$R_{t+1}^{-1}\left[\sum_{m=1}^{N}\pi_\Theta\left(\theta_{t+1}^m|\theta_t^n\right)\nu_m MC\left(\theta_{t+1}^m\right)-\sum_{m=1}^{N}\pi_\Theta\left(\theta_{t+1}^m|\theta_t^n\right)\left(\nu_{m+1}-\nu_m\right)DC\left(\theta_{t+1}^m\right)\right]$$

This object is equal to (minus) the right-hand side of (46), multiplied by R_{t+1}^{-1}. By Lemma 3 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 3 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 reduce the within-period utility of a mimicker by $\beta\delta$ units. By earlier arguments, the marginal cost of this perturbation as δ is moved away from zero, assessed at time t, will be:

$$\beta\pi_\Theta\left(\theta^t\right)DC\left(\theta_t^n\right)$$

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.

A.10 Proof of Lemma 5

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.

Consider the following construct. Reduce the utility of an agent of type θ_{t+1}^1 (with the relevant prior history) by a unit *through a reduction in consumption alone*. This will necessarily reduce the utility of an agent of type θ_{t+1}^2 who mimicks θ_{t+1}^1, so to preserve exact incentive compatibility at $t+1$ we must reduce the truth-telling utility of θ^2_{t+1} by a compensating amount. Suppose this is likewise done by reducing the consumption of that agent alone. Further reductions in utility must then be provided to θ^3_{t+1}, again assumed to be done through consumption changes alone, and so on up to θ^N_{t+1}. The consumption of all agents will have fallen at $t+1$, and therefore their utility will have done likewise. Suppese that the expected reduction in $t+1$ utility, assessed in period t, is some amount $\overline{\nu}$. Then incentive compatibility can be preserved from the perspective of period t by raising the within-period utility of the relevant type θ_t by an amount $\beta\overline{\nu}$, along a dimension in consumption-output space that has an equal effect on 'one-higher' minickers. But since $\alpha(\theta_t) < 1$ the marginal cost of this period-t perturbation will be positive, whilst by construction the consumption changes at $t+1$ generate positive resources. Hence the combined perturbation generates a surplus, contradicting optimality.[51]

A.11 Proof of Proposition 11

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 6 $\frac{\tau(\theta^n_t)}{u_c\left(\widehat{\theta}^n_t;\theta^{n+1}_t\right)(1-\tau(\theta^n_t))+u_y\left(\widehat{\theta}^n_t;\theta^{n+1}_t\right)} \overset{a.s.}{\to} 0$ *holds under an optimal plan that solves the restricted problem.*

Proof. In the iid case this follows directly from equation (33):

$$\begin{aligned}
&\lim_{t\to\infty}\left[-\pi_\Theta\left(\theta^n_{t+1}|\theta_t\right)\frac{\tau\left(\theta^n_{t+1}\right)}{u_c\left(\widehat{\theta}^n_{t+1};\theta^{n+1}_{t+1}\right)\left(1-\tau\left(\theta^n_{t+1}\right)\right)+u_y\left(\widehat{\theta}^n_{t+1};\theta^{n+1}_{t+1}\right)}\right] \qquad (84)\\
=&\ -\sum_{m=n+1}^{N}\pi_\Theta\left(\theta^m_{t+1}|\theta_t\right)\lim_{t\to\infty}\left[\frac{1-\alpha\left(\theta^m_{t+1}\right)}{u_c\left(\theta^m_{t+1}\right)+u_y\left(\theta^m_{t+1}\right)\alpha\left(\theta^m_{t+1}\right)}\right]\\
&+\pi_\Theta\left(\theta_{t+1}>\theta^n_{t+1}|\theta_t\right)\sum_{\theta_{t+1}\in\Theta}\pi_\Theta\left(\theta_{t+1}|\theta_t\right)\lim_{t\to\infty}\left[\frac{1-\alpha\left(\theta_{t+1}\right)}{u_c\left(\theta_{t+1}\right)+u_y\left(\theta_{t+1}\right)\alpha\left(\theta_{t+1}\right)}\right]\\
=&\ 0
\end{aligned}$$

In the Markov case we know that equation (33) 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 $\theta=\theta^N$ has obtained along

[51] Notice that this argument cannot be applied in the case of non-iid productivity processes, since the perturbation 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.

the given sample path, we must have:

$$
\begin{aligned}
& \lim_{T\to\infty}\left[-\pi_\Theta\left(\theta^n_{t(T)+1}|\theta_{t(T)}\right)\right. \qquad (85)\\
& \left.\cdot\frac{\tau\left(\theta^n_{t(T)+1}\right)}{u_c\left(\widehat{\theta}^n_{t(T)+1};\theta^{n+1}_{t(T)+1}\right)\left(1-\tau\left(\theta^n_{t(T)+1}\right)\right)+u_y\left(\widehat{\theta}^n_{t(T)+1};\theta^{n+1}_{t(T)+1}\right)}\right]\\
= & -\sum_{m=n+1}^{N}\pi_\Theta\left(\theta^m_{t(T)+1}|\theta_{t(T)}\right)\lim_{T\to\infty}\left[\frac{1-\alpha\left(\theta^m_{t(T)+1}\right)}{u_c\left(\theta^m_{t(T)+1}\right)+u_y\left(\theta^m_{t(T)+1}\right)\alpha\left(\theta^m_{t(T)+1}\right)}\right]\\
& +\pi_\Theta\left(\theta_{t(T)+1}>\theta^n_{t(T)+1}|\theta_{t(T)}\right)\\
& \cdot\sum_{\theta_{t(T)+1}\in\Theta}\pi_\Theta\left(\theta_{t(T)+1}|\theta_{t(T)}\right)\lim_{T\to\infty}\left[\frac{1-\alpha\left(\theta_{t(T)+1}\right)}{u_c\left(\theta_{t(T)+1}\right)+u_y\left(\theta_{t(T)+1}\right)\alpha\left(\theta_{t(T)+1}\right)}\right]\\
= & \; 0
\end{aligned}
$$

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 (46) and (45). 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 (54) we have:

$$
\begin{aligned}
u_c\left(\theta^n_t\right)+u_y\left(\theta^n_t\right)\alpha\left(\theta^n_t\right) &= \frac{u_c\left(\widehat{\theta}^n_t;\theta^{n+1}_t\right)-u_y\left(\widehat{\theta}^n_t;\theta^{n+1}_t\right)\frac{u_c\left(\theta^n_t\right)}{u_y\left(\theta^n_t\right)}}{1-\frac{u_y\left(\widehat{\theta}^n_t;\theta^{n+1}_t\right)}{u_y\left(\theta^n_t\right)}} \qquad (86)\\
&> u_c\left(\widehat{\theta}^n_t;\theta^{n+1}_t\right)-u_y\left(\widehat{\theta}^n_t;\theta^{n+1}_t\right)\frac{u_c\left(\theta^n_t\right)}{u_y\left(\theta^n_t\right)}\\
&= u_c\left(\widehat{\theta}^n_t;\theta^{n+1}_t\right)-u_y\left(\widehat{\theta}^n_t;\theta^{n+1}_t\right)\frac{1}{\left(1-\tau\left(\theta^n_t\right)\right)}
\end{aligned}
$$

If

$$
u_c\left(\widehat{\theta}^n_t;\theta^{n+1}_t\right)\left(1-\tau\left(\theta^n_t\right)\right)+u_y\left(\widehat{\theta}^n_t;\theta^{n+1}_t\right)\overset{a.s.}{\to}\infty
$$

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 $(1-\tau(\theta_t^n))\in[0,1]$ follows from the definition of τ and Proposition 7. 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 $\frac{1-\alpha(\theta_t^n)}{u_c(\theta_t^n)+u_y(\theta_t^n)\alpha(\theta_t^n)}$ can only converge to a non-zero limit if $|\alpha(\theta_t)|$ is itself always infinite at that limit. But since we know $\alpha(\theta_t)=0$ when $\theta_t=\theta^N$ we can rule that out.

The alternative is that $\tau(\theta_t^n)\overset{a.s.}{\to}0$. In this case we have $u_c(\theta_t^n)=-u_y(\theta_t^n)$ 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 $u_c(\theta_t^n)=-u_y(\theta_t^n)$ 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.