Review of Economic Studies (2024) **91**, 599–640 https://doi.org/10.1093/restud/rdad043

Advance access publication 6 April 2023

Equilibrium Analysis in Behavioural One-Sector Growth Models

Daron Acemoglu
Department of Economics, Massachusetts Institute of Technology

and

Martin Kaae Jensen
School of Economics, University of Nottingham

First version received December 2020*; Editorial decision December* 2022*; Accepted April* 2023 *(Eds.)*

Rich behavioural biases, mistakes, and limits on rational decision-making are often thought to make equilibrium analysis much more intractable. We establish that this is not the case in the context of one-sector growth models such as Ramsey–Cass–Koopmans or Bewley–Aiyagari models. We break down the response of the economy to a change in the environment or policy into two parts: the direct response at the given (pre-tax) prices, and the equilibrium response which plays out as prices change. Our main result demonstrates that under weak regularity conditions, regardless of the details of behavioural preferences, mistakes and constraints on decision-making, the long-run equilibrium will involve a greater capital-labour ratio if and only if the direct response (from the corresponding consumption-saving model) involves an increase in aggregate savings. One implication of this result is that, from a qualitative point of view, behavioural biases matter for long-run equilibrium if and only if they change the direction of the direct response. We provide detailed illustrations of how this result can be applied and generate new insights using models of misperceptions, self-control and temptation, and naive and sophisticated quasi-hyperbolic discounting.

Key words: Behavioural economics, Comparative statics, General equilibrium, Neoclassical growth

JEL codes: D90, D50, O41

1. INTRODUCTION

Most standard macro and growth models rely on very restrictive behavioural assumptions about households—infinitely lived, often representative, agents that are capable of solving complex maximization problems without any behavioural biases or limitations, and of implementing the optimal decisions without any inconsistencies or mistakes. It is an uncomfortable stage of introductory graduate courses when these assumptions are introduced and students rightfully ask whether everything depends on them. A natural conjecture is that these assumptions do matter and any degree of behavioural richness would render any general conclusions impossible. Not only do general equilibrium effects become notoriously complicated and the set of indirect

The editor in charge of this paper was Kurt Mitman.

effects correspondingly rich; we would also expect the specific departure from full rationality—*e.g.* systematic mistakes, ambiguous beliefs, overoptimism, or dynamic inconsistency—to have a first-order impact on the direction in which the economy responds to changes in policy or technology.

In this paper, we study one-sector growth models and establish that while it is true that at the individual level outcomes depend critically on the exact behavioural specification, robust predictions of long-run responses to changes in environment (policy, preferences or technology) can nonetheless be obtained in the presence of general behavioural preferences. Specifically, we identify conditions that are sufficient—and when the steady-state equilibrium is unique or when changes are small, also necessary—for changes in environment to lead to comparative statics in line with the predictions of the baseline neoclassical growth models. These conditions depend only on the direction of the *direct response* to a change in environment, defined as the (partial equilibrium) impact on aggregate savings, computed from the consumption-saving problem of households, holding the pre-tax prices fixed at their initial steady-state values. Put simply, if the direct response to a change in environment is an increase in aggregate savings, then no matter how complex the general equilibrium interactions that will play out dynamically (as prices change), the long-run impact on the capital stock and output per capita will be positive. Conversely, if the direct response is a decrease in aggregate savings, then the long-run impact on the capital stock and output per capita will be negative.

Before we elaborate on this result and provide an intuition, let us explain it in the context of a specific policy change—a reduction in the capital income tax rate. In baseline "neoclassical" settings, including the Ramsey–Cass–Koopmans model or the Bewley–Aiyagari model, the direct response is simply the "partial equilibrium" change in aggregate savings, holding prices at their initial steady-state values. This direct response is positive under standard assumptions, and in this case, so is the long-run response: lower capital income taxes lead to higher capital-output ratio and output per capita in the long run. Taking this as a benchmark, our results can then be read as saying that *any* set of rich and more realistic behavioural preferences that do not reverse the direction of the partial-equilibrium response leave the *qualitative* comparative statics of the steady-state equilibrium unchanged—the capital-labour ratio and output per capita will increase following a reduction in capital income taxes.[1] These results apply with minimal assumptions and allow households to have different behavioural preferences and make various systematic mistakes.

Conversely, our results also delineate robust conditions for behavioural preferences and systematic mistakes to *reverse* the direction of long-run comparative statics: when the direct response to a change in environment goes in the opposite direction of the direct response in benchmark neoclassical models, long-run (general equilibrium) comparative statics will also go in the opposite direction of the conventional comparative statics. So if lower capital taxes reduce aggregate savings upon impact, they will lead to lower capital stock and output per capita in the long run.

Figure 1 presents these results diagrammatically. All four panels of the figure depict a key object in our analysis, "*the market correspondence*," which summarizes the aggregate savings responses at different levels of the capital-labour ratio (see Section 2.5). Our main theorem amounts to saying that, for long-run comparative statics, it is sufficient to look at how the market correspondence shifts at the capital-labour ratio of the initial steady-state equilibrium. Panel A illustrates this point. Even though the market correspondence that applies for a new environment

1. Naturally, different distribution of preferences and mistakes across households will have *quantitative* implications. These are of course important for many applications, even though they are not our focus in the current paper.

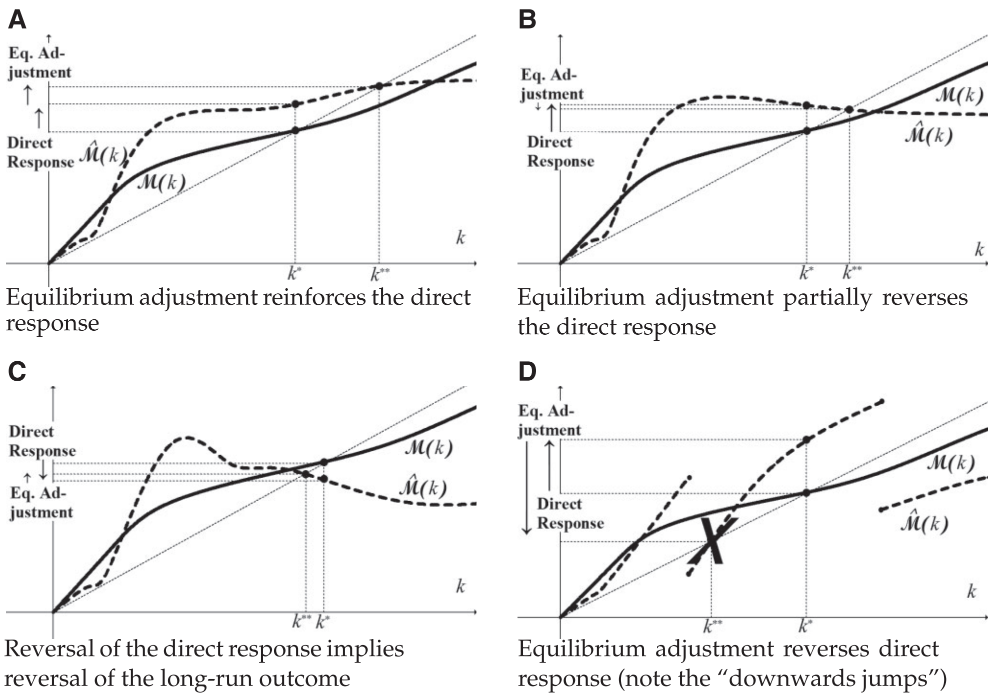

FIGURE 1

Panel A shows an instance in which general equilibrium effects amplify the direct response, while in Panel B they dampen it. In Panel C, the direct response is a decline in aggregate savings, so the long-run impact incorporating general equilibrium effects is also negative. The scenario in Panel D, where the direct response is positive and the long-run impact is negative, is impossible in the one-sector behavioural growth model because individual saving functions cannot "jump down" (equivalently, consumption functions cannot "jump up"). To overturn the (long-run) comparative statics in Panels A–B, the direct response must be negative as in Panel C

is not everywhere above the initial market correspondence, it is strictly above it at the original capital-labour ratio, and this is sufficient for us to establish that the change in environment will lead to a higher capital-labour ratio.

Panel B provides a complementary configuration. While in Panel A general equilibrium interactions reinforced the direct response, in this case they dampen it. In general, it is very difficult to determine, without explicit computations, whether Panel A or Panel B will apply—because general equilibrium interactions are difficult to characterize. Crucially, however, the direction of long-run comparative statics can be determined without this knowledge.

Panel C depicts the converse case. Now the direct response is a reduction in aggregate savings. As a result, the figure shows that the long-run and output per capita will decline. Hence, if we think of Panel A as corresponding to the benchmark neoclassical growth model, Panel C represents the case where behavioural preferences reverse the direction of the direct response, and thus lead to the complete opposite of the neoclassical long-run comparative statics. Finally, Panel D depicts the case ruled out by our theorems. The configuration in this last panel illustrates that, in principle, there is nothing automatic about our results (in fact, this is particularly the case once we are in the case with more than a single aggregate). Nevertheless, we will show that this configuration cannot arise when there are no downward jumps in the market correspondence, which can be guaranteed under fairly weak assumptions in the one-sector model.

To build intuition for our results, let us first revisit the standard Bewley–Aiyagari model with fully rational heterogeneous agents. In such an economy, the equilibrium adjustment following the direct response involves random/stochastic changes in the distribution of assets, as well as prices and the aggregate capital stock as the economy settles into a new steady-state equilibrium. Even with fully rational agents, this adjustment is complex: because of income effects, some households may change their savings in the opposite direction of the aggregate change as their income and the prices they face evolve. With behavioural preferences or biases, it is potentially even more so, since we have to take into account not just the conventional income effects and price changes, but also any systematic mistakes in optimization or expectations, more complex intertemporal trade-offs and issues related to dynamic inconsistency. Nonetheless, our main theorems show that, even in such settings and exactly at the same level of generality as in the baseline Bewley–Aiyagari economy, we can establish qualitative long-run comparative statics. Our analysis also establishes that although fairly general results about aggregate changes can be derived, there is a type of "indeterminacy" at the individual level—nothing much can be said about how individuals will behave and which individuals will go in the opposite direction of the aggregate economy. This observation further clarifies that our results are not a consequence of some (implicit) monotonicity assumption that ensures all households move in the same direction. On the contrary, our results are about aggregate outcomes, without any knowledge or implications on how any given household will adjust.

We can now present the intuition for these results at two complementary levels. The first is economic in nature and it is related to an idea that already appears in Becker (1962) that "aggregation" disciplines economic behaviour. Though we cannot say anything about individual behaviour, we can determine the behaviour of market-level variables (*i.e.* aggregates such as the capital stock and income per capita). This is because even if many households respond in the opposite direction of the direct response, in equilibrium enough households have to move in the same direction as the direct response.

The second intuition for our result is more mathematical. Suppose that the steady-state equilibrium is unique, and focus on a policy change that increases aggregate savings at the initial capital-labour ratio. Then the only way the new steady-state equilibrium could have lower capital stock is when the equilibrium response goes in the opposite direction and more than offsets the initial increase in aggregate savings. This in turn can only be true if a higher capital stock induces lower savings. But even if this were the case, the equilibrium response could not possibly overturn the direct response. This is because the economic force leading to lower savings would not be present if the new steady-state equilibrium ended up with a lower capital stock, and thus the indirect equilibrium response could not overturn the initial (positive) direct response. When there are multiple steady-state equilibria, this reasoning would not apply to all of them, but we develop a similar argument for extremal (greatest and least) steady-state equilibria.

In Section 4, we use several popular behavioural models to further illustrate our theorems and show how they can be applied fairly straightforwardly, yielding new insights. We start with a model of persistent misperceptions and establish how the form of misperceptions matters and leads to different types of results, and also demonstrate how they can sometimes reverse standard comparative statics (*e.g.* lower taxes on capital income reducing the long-run capital-labour ratio). We then show how our main theorems lead to new comparative statics in the context of macro models incorporating self-control and temptation problems as in Gul and Pesendorfer (2004). Finally, we discuss naive and sophisticated versions of quasi-hyperbolic preferences introduced and analysed in Strotz (1956), Phelps and Pollak (1968), Laibson (1997) and Harris and Laibson (2001). In this context, we also show how our theoretical results can be blended with simple numerical analysis.

Our paper is related to several literatures. The first, already mentioned, is Becker's seminal paper which argues that market demand curves will be downward sloping even if households are not rational because their budget constraints will put pressure for even random behaviour to lead to lower demand for goods that have become more expensive. Machina (1982) makes a related observation about the independence axiom in expected utility theory. Though related to and inspired by these contributions, our main result is very different. While Becker's argument is about whether an increase in price will lead to a (partial equilibrium) change in aggregate behaviour consistent with "rational behaviour," our focus is about taking the initial change in behaviour, whether or not it is rational, as given and then establishing that, under general conditions on the objectives and behavioural biases and constraints of households, the (general) equilibrium responses will not reverse this direct response.

The second literature we build on is robust comparative statics (*e.g.* Topkis, 1978; Vives, 1990; Milgrom, 1994; Milgrom and Roberts, 1994; Milgrom and Shannon, 1994; Quah, 2007). Not only do we share these papers' focus on obtaining robust qualitative comparative static results, but we also use similar tools, in particular a version of the "curve-shifting" arguments of Milgrom and Roberts (1994) (see also Acemoglu and Jensen, 2015) which allow us to derive robust results in non-monotone economies.[2] Nevertheless, our main theorem is not an application of any result we are aware of. Rather, it significantly extends and strengthens the approach used in the robust comparative statics literature (we provide a detailed technical discussion of the relationship with previous literature in Appendix B). Most significantly, in contrast to other approaches in the literature, our comparative static results only rely on "local information"—on behaviour at a *specific* capital-labour ratio (or vector of prices) rather than the much stronger notions requiring that behaviour increases or decreases savings for all prices.[3] As a result, we are able to establish economically and mathematically stronger results: whenever the sum of the initial savings responses of all agents is positive at the initial capital-labour ratio, the full general equilibrium will involve an increase in the capital-labour ratio.

In this context, it is also useful to compare our results to those of our earlier paper, Acemoglu and Jensen (2015), where we analysed a related setup, but with three crucial differences. First, and most importantly, there we focused on rational households, thus eschewing any analysis of behavioural biases and their impacts on equilibrium responses. Second, and as a result of the first difference, we did not have to deal with the more general problem considered here, which requires a different mathematical approach. Third and also crucially, we imposed considerably stronger assumptions to ensure that the direct response of all households went in the same direction at *all* prices, which we do not do in the current paper.[4]

Finally, our paper is related to several recent works that incorporate rich behavioural biases and constraints into macro models. In addition to those already mentioned, these include Krusell and Smith (2003), Krusell *et al.* (2010), and Cao and Werning (2018) who study the dynamic and equilibrium implications of hyperbolic discounting;[5] temptation and self-control preferences

2. See p. 590 in Acemoglu and Jensen (2015) for additional discussion of such non-monotone equilibrium comparative statics results.

3. See for example Lemma 1 (and Figures 1–3) in Milgrom and Roberts (1994) or Definition 5 in Acemoglu and Jensen (2015). Milgrom and Roberts (1994) also use local assumptions, but just to derive local comparative statics results (see Figure 7 and the surrounding discussion); this is different from our results, which are global despite being based on local assumptions.

4. An alternative and complementary approach is based on mean-field games. Particularly noteworthy is Light and Weintraub (2021), who investigate comparative statics in mean-field games, but once again focusing on uniform and global changes (see, *e.g.* their Theorem 4). Ahn *et al.* (2018) and Achdou *et al.* (2022) apply mean-field game techniques to the Bewley–Aiyagari model.

5. Barro (1999) shows that, with logarithmic utility and a representative household, hyperbolic discounting leads to similar insights to the standard one-sector growth models. His results do not extend beyond the logarithmic case and representative household models, and are not related to our general approach.

as in Gul and Pesendorfer (2001, 2004) and Fudenberg and Levine (2006, 2012); non-separable preferences in dynamic macro models as in Koopmans (1960), Epstein and Hynes (1983), Kreps and Porteus (1978), Epstein and Zin (1989, 1991), and Backus *et al.* (2004); models of ambiguity and multiple priors as in Gilboa (1987) and Gilboa and Schmeidler (1995); and models of sparse optimization as in Gabaix (2014, 2017).

The rest of the paper is organized as follows. Section 2 describes the model and introduces the "market correspondence" (which is key to our analysis). Section 3 contains the main results and applications. Section 4 shows how our results can be applied in the presence of systematic misperceptions, self-control and temptation preferences, and quasi-hyperbolic households. Section 5 concludes, Appendixes A and B contain the proofs of most of the results stated in the text and additional results, with the remaining proofs presented in the online Appendix C (supplementary material).

2. BEHAVIOURAL ONE-SECTOR GROWTH MODELS

This section introduces our general setup, which blends a standard growth model with various behavioural preferences.

2.1. *Production and markets*

The production side is the same as the canonical neoclassical growth model (*e.g.* Acemoglu, 2009) augmented with general distortions.

Labour is in fixed supply and normalized to unity so we can use capital, capital-labour ratio and capital-per-worker interchangeably and denote it by k. Markets clear at all times, and production is described by a profit maximizing aggregate constant returns firm with a smooth (per capita) production technology $y = f(k)$ that satisfies $f(0) = 0$, $f' > 0$, and $f'' < 0$. We also impose that there exists $\bar{k} > 0$ such that $f(k) < k$ all $k \geq \bar{k}$, which ensures compactness. This condition is implied by the standard Inada conditions when these are imposed. The rate of depreciation is $\Delta \in [0, 1]$.

We allow for taxes and distortions $\omega(k)$ and $\tau(k)$ on labour and capital. Throughout, "market prices" refer to pre-tax factor prices, $\hat{w}(k_t) \equiv f(k_t) - f'(k_t)k_t$ and $\hat{R}(k_t) \equiv f'(k_t)$. Hence, the after-tax (and after-distortion) wage and rate of return facing the households are

$$w_t = w(k_t) \equiv (1 - \omega(k_t))(f(k_t) - f'(k_t)k_t) \tag{1}$$

and

$$R_t = R(k_t) \equiv (1 - \tau(k_t))f'(k_t) - \Delta. \tag{2}$$

The simplest example of such distortions are proportional taxes on capital and labour income, $\tau(k_t) = \tau$ and $\omega(k_t) = \omega$. Other examples include distortions from contracting frictions or markups due to imperfect competition. When $\tau(k) = \omega(k) = 0$ for all k, we recover the benchmark case with no distortions.

We allow proceeds from these distortions to be partially rebated to households (which will be the case when they represent taxes and some of the tax revenues are redistributed to households or when they result from markups that generate profits).

The total amount of resources that is *not* rebated to households—that is, either consumed by the government, invested in public goods or wasted, in all cases in a way that does not affect marginal utilities—is denoted by $G(k_t)$. If nothing is rebated, then $G(k_t) = \omega(k_t)(f(k_t) - f'(k_t)k_t) + \tau(k_t)f'(k_t)$. On the other hand, if the only source of distortions is taxes because the government rebates everything back to consumers (*e.g.* in the form of lump-sum transfers), then $G(k_t) = 0$.

2.2. *Households and capital markets*

There is a continuum of households [0, 1] with a typical household denoted by $i \in [0, 1]$. As in Aiyagari (1994), households are subject to borrowing constraints and can save either by investing in a riskless government bond that is in zero net supply or in the capital stock of the economy, k_t. Throughout, we assume that any randomness is such that there is no aggregate uncertainty, which ensures that capital k_t is deterministic and factor prices are given by (1) and (2) at all times.[6]

At date t each household i is subject to a labour endowment shock denoted by $l_t^i \in [l_{\min}^i, l_{\max}^i] \subset \mathbb{R}_{++}$ and a preference shock $\epsilon_t^i \in E^i \subseteq \mathbb{R}$ (where we take E^i to be compact). We assume that $(l_t^i, \epsilon_t^i)_{t=0}^{\infty}$ follows a Markov process with invariant distribution μ_i. It is convenient to set $e_t^i = (l_t^i, \epsilon_t^i, w_t, R_t, T_t^i)$, where T_t^i denotes the transfers/rebates that household i receives at time t.

Household i's objective is to maximize utility conditional on its beliefs (or expectations) about the future variables $(e_\tau^i)_{\tau=t+1}^{\infty}$ as well as its anticipated future savings behaviour. Let us denote the *true model* by θ^M. This includes a complete description of all of this section's contents, including current and future taxes, the stochastic process governing $(l_t^i, \epsilon_t^i)_{i\in[0,1]}$, equilibrium conditions, and so forth.

Household $i \in [0, 1]$ forms beliefs at date t on the basis of the true model θ^M and its observations of economic variables summarized in e_t^i. We suppress the dependence on the true model θ^M throughout to reduce notation, and summarize the belief process with the mapping $P_t^i : e_t^i \mapsto P_t^i(\cdot; e_t^i)$, which defines a probability measure on future outcomes given the current vector of variables e_t. That is, for any (Borel) measurable set of future observations B, the household believes that $(e_\tau^i)_{\tau=t+1}^{\infty}$ will lie in B with probability $P_t^i(B; e_t^i) \in [0, 1]$.[7] Rational expectations is the special case of this formulation, where the marginal distribution of exogenous parameters coincides with objective probabilities implied by the Markov process $(l_t^i, \epsilon_t^i)_{\tau=t+1}^{\infty}$, and the household uses the true model θ^M to correctly predict future prices. A simple and familiar example is the Bewley–Aiyagari model (Aiyagari, 1994) with *i.i.d.* labour endowment shocks $l^i \sim \mu_i$. Because in this case agents have rational expectations, beliefs about future prices coincide with actual (equilibrium) prices and beliefs about the future realizations of the labour endowment shock coincide with the objective probability measure, μ_i. For this reason, as in models with rational expectations more generally, beliefs can be suppressed/ignored altogether.

Other belief formation processes may completely ignore the true model and specify "dogmatic" misperceptions that are not revised (even when they contradict the data repeatedly), or generate beliefs on the basis of other variables summarized in e_t^i, which may involve some Bayesian or non-Bayesian updating. In particular, unlike in models based on rational expectations and common knowledge, households' beliefs may be in contradiction with each other and

6. Like in Aiyagari, a riskless arbitrage condition ensures that households are indifferent between investing in government bonds and the capital stock. If government bonds are in positive net supply as in Aiyagari and McGrattan (1998), then the analysis needs to be modified along the lines of Aiyagari and McGrattan (1998, pp. 452–453), but their arguments establish that this is still a one-sector economy and thus all of our results apply.

7. Formally, $P_t^i(\cdot; e_t^i)$ is a regular Borel (probability) measure on the set of future observations, which can be taken to be the space of bounded infinite sequences with the supremum norm. For technical reasons, we restrict attention to measures whose conditional probability of l_{t+1}^i lying in a measurable subset of $\mathbb{R}_+$ has continuous Radon–Nikodym derivative with respect to the Lebesgue measure on $\mathbb{R}_+$.

with actual outcomes. Note also that since w_τ and R_τ are the after-tax/after-distortion wage and rate of return, P_t^i implicitly incorporates (potentially incorrect) beliefs about future taxes and distortions.

If at date t, household i consumes c_t^i and its future consumption is $(c_{t+1}^i, c_{t+2}^i, \ldots)$, then it gets utility

$$u^{\epsilon_t^i}(c_t^i) + V^{i,\epsilon^{i,t+1}}(c_{t+1}^i, c_{t+2}^i, \ldots). \tag{3}$$

Here $u^{\epsilon_t^i}$ and $V^{i,\epsilon^{i,t+1}}$ are the current and future utility functions, respectively, and $(\epsilon_t^i, \epsilon^{i,t+1}) = (\epsilon_t^i, \epsilon_{t+1}^i, \epsilon_{t+2}^i, \ldots)$ is a sequence of parameters that directly impact utility. These parameters could reflect discount factors or idiosyncratic tastes or biases as in McFadden (1974, p. 108).[8] If $u^{\epsilon_t^i,i} = u^i$ and

$$V^{i,\epsilon^{i,t+1}}(c_{t+1}^i, c_{t+2}^i, \ldots) = \sum_{\tau=t+1}^{\infty} \beta\delta^{\tau-t} u^i(c_\tau^i), \quad \delta < 1, \quad \beta > 0$$

we get the time-separable, geometrically discounted benchmark case when $\beta = 1$, and the quasi-hyperbolic model when $\beta \neq 1$. A third example is when the household has a finite time-horizon, $V^{i,\epsilon^{i,t+1}}(c_{t+1}^i, c_{t+2}^i, \ldots) = \sum_{\tau=t+1}^{t+T} \delta^{\tau-t} u^i(c_\tau^i)$. This may be interpreted as a simple version of "sparsity" in the sense of Gabaix (2014, 2017).[9]

We denote household i's assets by a_t^i and its borrowing constraint by $\underline{a}^i \leq 0$ (which is assumed to be above the solvency constraint; see *e.g.* Aiyagari, 1994, p. 666). We also impose an upper bound $\overline{a}^i > \underline{a}^i$, but this comes with no loss of generality under compactness in production (Section 2.1) as $\overline{a}^i$ may be chosen so that it never binds in equilibrium, P_t^i-almost everywhere and for almost every household i.

If at date t, the household chooses (gross) savings $s_t^i = a_{t+1}^i \in [\underline{a}^i, \overline{a}^i]$, its current consumption will be

$$c_t^i = (1 + R_t)a_t^i + w_t l_t^i + T_t^i - a_{t+1}^i, \tag{4}$$

and, conditioning on $(e_\tau^i)_{\tau=t+1}^{\infty}$, its future consumption will be

$$\begin{aligned} c_{t+1}^i &= (1 + R_{t+1})a_{t+1}^i + w_{t+1}l_{t+1}^i + T_{t+1}^i - \tilde{s}_{t+1}^i(a_{t+1}^i; e_{t+1}^i) \\ c_{t+2}^i &= (1 + R_{t+2})\tilde{s}_{t+1}^i(a_{t+1}^i; e_{t+1}^i) + w_{t+2}l_{t+2}^i + T_{t+2}^i - \tilde{s}_{t+2}^i(\tilde{s}_{t+1}^i(a_{t+1}^i; e_{t+1}^i); e_{t+2}^i) \\ \vdots &= \qquad \vdots \end{aligned} \tag{5}$$

where $s_\tau^i(a_\tau^i; e_\tau^i)$ denotes anticipated savings of a "future self" at date $\tau > t$ conditioned on future asset a_τ^i and e_τ^i. In this formulation, we are taking future saving functions, s_τ^i for $\tau > t$, as given. When the household is dynamically consistent and fully rational, she chooses them as her policy functions. In general, they may be given by the choices of future selves with only imperfectly aligned interests, and/or misperceived. Note also an important convention: Throughout,

8. Models such as Gul and Pesendorfer (2004) where utility also depends on asset holdings (Section 4.3) are nested in this formulation, since asset levels are a function of past consumption choices given initial asset holdings and factor prices.

9. In the working paper version, we also studied the recursive utility specification of Epstein and Zin (1989), which relaxes additivity in (3), and we further extend what follows to non-additive beliefs and Choquet expected utility (ambiguity aversion).

we assume that the functions s^i_τ for $\tau > t$ are always correctly perceived, and any misperceptions about the behaviour of future selves are represented via the utility parameter ϵ_τ.[10]

At date t, a (gross) savings level $a^{i,*}_{t+1}$ is *optimal* given current assets a^i_t, current observations e^i_t, the (measurable) anticipated future saving functions s^i_τ, $\tau > t$ and the beliefs $P^i_t(\cdot; e^i_t)$ about future variables $(e^i_\tau)^\infty_{\tau=t+1}$ if it maximizes the expected value of (3) subject to (4) and (5). Denoting a sequence of future variables by $\tilde{e}$ and substituting for consumption, we can write this compactly as

$$a^{i,*}_{t+1} \in \arg \max_{a' \in [\underline{a}^i, \min\{y^i_t, \bar{a}^i\}]} u^{i,\epsilon^i_t}(y^i_t - a') + \int W^i(a', \tilde{e}; (s^i_\tau)^\infty_{\tau=t+1}) P^i_t(d\tilde{e}; e^i_t), \tag{6}$$

where $y^i_t = (1 + R_t)a^i_t + w_t l^i_t + T^i_t$ is wealth and the continuation utility W^i is a measurable function given as

$$W^i(a', (e^i_\tau)^\infty_{\tau=t+1}; (s^i_\tau)^\infty_{\tau=t+1}) = V^{i,\epsilon^{i,t+1}}((1 + R_{t+1})a' + w_{t+1}l^i_{t+1} + T^i_{t+1} - s^i_{t+1}(a'; e^i_{t+1}),$$
$$(1 + R_{t+2})s^i_{t+1}(a'; e^i_{t+1}) + w_{t+2}l^i_{t+2} + T^i_{t+2} - s^i_{t+2}(s^i_{t+1}(a'; e^i_{t+1}); e^i_{t+2}), \ldots).$$

Observe that the benchmark Bewley–Aiyagari model is a special case of this formulation. In this case, P^i_t coincides with the true marginal distribution of exogenous parameters and places probability 1 on the actual values of future endogenous variables $((w_t, R_t)^\infty_{\tau=t+1})$, given the Markov process for $(l^i_t)^\infty_{\tau=t+1}$; s^i_t is directly determined from the households' dynamic programming problem; and the continuation utility W^i can be obtained from standard dynamic programming. More generally, however, (6) nests various behavioural biases or dynamic inconsistencies such as when households have misperceptions about the future (Section 4.2), or when discounting is hyperbolic (Section 4.4).

In what follows, we assume that utility functions and beliefs are continuous. Because the set of feasible assets is uniformly bounded, this is sufficient to ensure a uniformly bounded objective function in (6). All sequence spaces are equipped with the supremum norm and the Borel σ-algebra, and the topology on probability measures is the weak convergence topology (*e.g.* see Epstein and Zin, 1989, p. 940).

Assumption 1. *u^i is a continuous, strictly increasing and strictly concave function, V^i is a continuous and strictly increasing function, and $P^i_t(\cdot; e^i)$ is continuous in e^i.*

2.3. *Time-stationary saving correspondences*

We say that beliefs for household i are *time-invariant* if for all $t = 1, 2, 3, \ldots$, we have

$$P^i_t = P^i \quad \text{for all } t = 0, 1, 2, \ldots. \tag{7}$$

The next definition imposes time-invariant belief processes and also requires that current selves expect future selves to adopt the same saving function.

Definition 1 (Time-stationary saving functions and correspondences)**.** s^i is a *time-stationary saving function* (*TSSF*) if for all initial levels of assets $a^i \in [\underline{a}^i, \bar{a}^i]$, all (w, R, T^i), and almost

10. One could derive from $P^i_t(\cdot; e^i_t)$ an induced probability measure over the space of saving functions, but working directly with $P^i_t(\cdot; e^i_t)$ is notational and conceptually simpler.

all z^i:

$$s^i(a^i;e^i) \in \arg \max_{a' \in [\underline{a}^i, \min\{y^i, \overline{a}^i\}]} u^{i,\epsilon_0^i}(y^i - a') + \int W^i(a', e'; s^i) P^i(de'; e^i), \tag{8}$$

where $e^i = (z^i, w, R, T^i)$, $y^i = (1+R)a^i + wl^i + T^i$, $W^i(a', e'; s^i) = W^i(a', e'; (s^i)_{\tau=t+1}^{\infty})$ and $e' = (e_0^i, e_1^i, \ldots)$. The union of all time-stationary saving functions is called the *time-stationary saving correspondence* , $S^i(a^i;e^i) = \{s^i(a^i;e^i) : s^i \text{ is a TSSF}\}$.

We emphasize that because Definition 1 allows beliefs to be incorrect, it nests both the case in which households are "sophisticated" (*e.g.* Strotz, 1956; Laibson, 1997; Harris and Laibson, 2001), and cases where agents are "naive" (in the sense of Strotz, 1956) and expectations are misaligned with future behaviour. It also nests recursive models such as Bewley–Aiyagari where s^i can be computed by standard dynamic programming.[11]

A correspondence is measurable if the inverse image of any open set is Borel-measurable (Aubin and Frankowska, 1990, p. 307). The proof of the next lemma is presented in Appendix A.

Lemma 1 (Basic properties of saving correspondences)**.** *Let Assumption 1 hold and suppose that each household's belief formation process is time-invariant. Then for each* $i \in [0,1]$*, the (time-stationary) saving correspondence* $S^i(a^i;e^i)$ *exists, is measurable in* (l^i, ϵ^i)*, upper hemi-continuous in* a^i*, w, R, and* T^i*, and its least and greatest selections are (weakly) increasing functions of assets* a^i*.*

They key observation is that under the general assumptions of the one-sector behavioural growth model summarized above, saving correspondences are "ascending" in the standard sense of robust comparative statics (*e.g.* Topkis, 1978; Vives, 1990; Milgrom and Roberts, 1994). This in particular means that the least and greatest selections (implied saving functions) from the saving correspondence are nondecreasing in assets. This is what rules out downward jumps in Figure 1 in the Introduction. An increasing saving correspondence implies that the associated least and greatest consumption functions increase less than one-for-one with assets. As a result, any consumption discontinuities must take the form of downward jumps—otherwise, there will be more than a one-for-one increase in consumption. Allowing for such discontinuities is important since these are common in the presence of dynamic inconsistencies (see *e.g.* Harris and Laibson, 2001, p. 937), as we will see explicitly in Section 4.5.

It is worth reiterating that by imposing time-invariance and focusing on time-stationary saving correspondences, we are greatly simplifying the description of the environment. First, time-invariance imposes time-stationary utility, so that households obtain the same continuation utility from the same consumption sequence starting from different points in time. Second, it ensures that the belief formation processes are time-invariant. One justification for time-invariance is that the environment may have already converged to a limit starting from some initial condition.

Time-invariance enables us to focus on the comparative statics of steady states, but is not without cost; our results have to be applied with care in settings that are not time-invariant.[12]

11. When beliefs are correct and discounting geometric, $\delta^{-1}W^i$ coincides with the standard value function obtained from dynamic programming.

12. For example, a policy change may create an initial period of belief confusion or mistaken perception, which becomes dissipated over time, inducing a specific type of time-dependence (Gabaix, 2017). If this is reversed in the course of the next $T < \infty$ periods, our steady-state analysis still applies in principle, but with some important caveats. This is because the relevant concept is no longer the "direct response" that takes place with the temporary beliefs, but the "hypothetical direct response" that would have obtained with the time-stationary beliefs (that apply after T periods) at the initial capital-labour ratio k^*.

2.4. *Steady-state equilibrium*

As a shorthand, from now on we define an "*environment*," denoted by $\theta = (\theta^M, (P^i)_{i\in[0,1]})$, to summarize the true model θ^M and the beliefs $(P^i)_{i\in[0,1]}$. Consider stationary market prices w and R and an environment θ (including stationary transfers T^i). Given these, λ^i is then an *invariant distribution for household i* if $\lambda^i(A\times B) = \int q^i(B;l^i,\epsilon^i)1_A(s^i(a^i;l^i,\epsilon^i,w,R,T^i))\lambda^i(da^i,d(l^i,\epsilon^i))$ where $q^i(B;l^i,\epsilon^i_t)$ is the *true* model's probability that $(l^i_{t+1},\epsilon^i_{t+1})$ lies in B given (l^i_t,ϵ^i_t), and s^i is a measurable selection from the time-stationary saving correspondence (here A and B are Borel subsets of $[\underline{a}^i,\overline{a}^i]$ and $[l^i_{\min},l^i_{\max}]\times E^i$, respectively).[13] The average (stationary) asset holding is then $\mathbb{E}[\hat{a}^i] = \int a^i\lambda^i(da^i,d(l^i,\epsilon^i))$ where $\hat{a}^i$ is the household's *stationary assets* given w, R, and θ (formally, $\hat{a}^i$ is the random variable on $[\underline{a}^i,\overline{a}^i]$ with distribution given by λ^i's marginal distribution of assets). In the Bewley–Aiyagari model, $\mathbb{E}[\hat{a}^i]$ is also the households' aggregate asset holdings.

Recall from (1) to (2) that w and R are the after-tax/distortions wage and rate of return, respectively. Hence they generally depend on the environment θ. Whenever this may cause confusion, we emphasize it by writing the market prices with the environment as a superscript. We define steady-state equilibria directly in terms of the corresponding capital-labour ratio, and also condition factor prices on the environment θ when this is necessary for emphasis or clarity.

Definition 2 (Equilibrium)**.** The capital-labour ratio $k^*\in\mathbb{R}_+$ represents a (*steady-state*) *equilibrium* given the environment θ, if equilibrium prices $w^* = w^\theta(k^*)$ and $R^* = R^\theta(k^*)$ are given by (1) *and* (2), household i's stationary asset distribution is $\hat{a}^{*,i}$ given w^*, R^*, and θ *for almost every* $i\in[0,1]$, and the capital market clears, that is, $k^* = \int \hat{a}^{*,i}\,di$.

Note that in this definition we are implicitly assuming that the households' aggregate asset holdings $\int \hat{a}^{*,i}\,di$ are well-defined by some version of the law of large numbers.[14] On the other hand, individual asset holdings will not be constant, though they will have a stationary distribution, which we denote by $\hat{a}^{*,i}$.

2.5. *The market correspondence*

We are now ready to formally define the key theoretical innovation of this paper, namely the *market correspondence.* We will see that steady states in our model correspond to intersections of the market correspondence with the 45° line (Lemma 2) and increases in (aggregate) savings translate into shifts in the market correspondence (Section 3).

Let $S = (S^i)_{i\in[0,1]}$ summarize the households' time-consistent savings correspondences. For a family of selections $s = (s^i)_{i\in[0,1]}\in S$, write $(\lambda^i(k;s))_{i\in[0,1]}$ if for almost every i, $\lambda^i(k;s) = \lambda^i$ is an invariant distribution when assets are scaled by $k/(\int \hat{a}^i(k;s)\,di)$, *i.e.* $\lambda^i(A\times B) = \int_{A\times l^i,\epsilon^i} q^i(B;l^i,\epsilon^i)1_A(s^i(a^i\frac{k}{\int \hat{a}^i\,di};l^i,\epsilon^i,w,R,T^i))\lambda^i(da^i,d(l^i,\epsilon^i))$, where $\int \hat{a}^i(k;s)\,di$ is

13. Since the least and greatest selections are increasing in assets (Lemma 1), there will exist an invariant distribution by Acemoglu and Jensen (2015), Theorem B1 and B3.

14. There is a large literature on laws of large numbers and their applications in continuum economies (*e.g.* Uhlig, 1996; Al-Najjar, 2004; Sun, 2006). Here and everywhere else in this paper we remain agnostic about precisely which formulation of the law of large numbers has been applied in the background. This "agnostic" approach is also the one taken in Acemoglu and Jensen (2015) where $\int \hat{a}^i(k)di$ is simply *assumed* to equal (or be one-to-one) with a real number. This approach has the advantage of not committing to a specific interpretation and therefore comes with maximum generality. On the downside, we must be careful to not push the generality of the setting too far: In the Aiyagari model, for example, any sensible application of a law of large numbers will require that the labour endowments' conditional distributions are at least pairwise independent conditioned on k. For further details and references, see Acemoglu and Jensen (2010, 2015).

the households' mean asset holdings. The implied distribution of consumption is denoted by $\hat{c}^i(k;s)$.[15] We then have:

Definition 3 (The Market Correspondence)**.** The *market correspondence* $\mathcal{M}^\theta : \mathbb{R} \to 2^{\mathbb{R}}$ *is:*[16]

$$\mathcal{M}^\theta(k) = \{f(k) + (1-\Delta)k - G(k)\} - \left\{ \int \hat{c}^i(k;s)di \in \mathbb{R}_+ : s \in S \right\}. \tag{9}$$

The right-hand side of the market correspondence (9) subtracts government and private sector consumption from total output plus unappreciated capital and thus gives the value of next period's capital stock. This motivates why, as in standard one-sector growth models, steady-state equilibria will be its fixed points.

The next lemma establishes that we can work directly with the market correspondence without specifying underlying equilibrium asset distribution. It also confirms that fixed points of the market correspondence will be steady-state equilibria. The proof of this lemma uses the fixed point comparative statics theorem of Acemoglu and Jensen (2015, Theorem 4, p. 601), which itself builds on Smithson's generalized fixed point theorem as well as Richter's theorem (Aumann, 1965). However, the most critical component of the proof is the observation that for a given k, $\mathcal{M}^\theta(k)$ equals the set of fixed points of a convex valued correspondence whose least and greatest selections are decreasing, and therefore it is itself convex-valued.

Lemma 2 (Properties of the market correspondence)**.** *Suppose that all households satisfy the assumptions in Lemma 1. Then the market correspondence $\mathcal{M}^\theta$ is a compact- and convex-valued upper hemi-continuous correspondence that begins above and ends below the 45° line. Furthermore, $k \in \mathcal{M}^\theta(k)$ if and only if k is a steady-state equilibrium.*

The market correspondence being convex-valued is an important and non-trivial property. This property does *not* follow from a convexification argument as in Aumann (1965), but depends critically on the fact that saving correspondences are increasing in the sense of Lemma 1 and so, in particular, on the fact that they have no jumps down. If, in fact, S^i were to have jumps down for a subset of agents of positive measure, then the correspondence $\int \mathcal{A}_k^{\theta,i}(\cdot)\, di$ in the proof would have jumps down as well. In that case, the market correspondence would not necessarily be convex-valued and this paper's main result that the average direct response determines the long-run outcome would become invalid.

3. MAIN RESULTS

This section contains our main results. Generalizations are provided in Appendix B and these results are applied in the context of specific behavioural models in Section 4.

Recall that θ^M denotes the "true model," $(P^i)_{i\in[0,1]}$ denotes the households' beliefs, and that the environment $\theta = (\theta^M, (P^i)_{i\in[0,1]})$ therefore contains all of the exogenous variables, parameters and policy variables of the model as well as specifications of how beliefs about

15. Precisely, $\hat{a}^i(k;s)$ has distribution equal to the marginal distribution of assets implied by $\lambda^i(k;s)$. $\hat{c}^i(k;s)$ has distribution $\lambda^i(k,s)(\{(a^i,l^i,\epsilon^i) : (1+R)a^i \frac{k}{\int \hat{a}^i\, di} + wl^i + T^i - s^i(a\frac{k}{\int \hat{a}^i\, di}; l^i, \epsilon^i, w, R, T^i) \in A\}$, where A is a Borel subset of the consumption set $\mathbb{R}_+$. $\hat{a}^i(k;s)$ and $\hat{c}^i(k;s)$ are well-defined under the assumptions of Lemma 1 (see the proof of Lemma 2).

16. This definition requires that the integral $\int \hat{c}^i(k;s)\, di$ has a degenerate distribution, and equation (9) refers to its (unit) mass point. Since $\hat{c}^i = (1+R(k))\hat{a}^i + w(k)\hat{l}^i + T^i - \hat{b}^i$, where $\hat{b}^i$ is the distribution of the next period's assets, this integral is well-defined whenever a law of large numbers applies (see footnote 14).

exogenous or endogenous objects are formed. This section studies changes in the environment and the *set of possible environments* Θ is taken to be an ordered set to facilitate this perspective.

For a given environment, $\theta^* \in \Theta$, say, we know from Lemma 2 that steady-state equilibria (Definition 2) correspond to points where the market correspondence intersects with the 45°-line, *i.e.* k^* is a steady state if and only if $k^* \in \mathcal{M}^{\theta^*}(k^*)$. This was illustrated in Figure 1 in the Introduction in the case where the market correspondence is single-valued (or we consider an appropriate selection from it).

We are now ready to define (individual and aggregate) direct responses discussed in the Introduction. In what follows, when this is necessary for emphasis, we condition the saving correspondence, as well as factor prices, on the environment θ.

Definition 4 (Individual direct responses)**.** Let k^* be an equilibrium given the environment $\theta^* \in \Theta$, and denote by $\lambda^{*,i}$ household i's associated invariant distribution. Let $\theta^{**} \in \Theta$ be a different environment. Then we say that *household i's direct response is positive at* k^* if its *asset holdings increase* at k^* when the environment changes from θ^* to θ^{**}, *i.e.* if

$$S^{\theta^{**},i}(a^i; e^{**,i}) \geq S^{\theta^*,i}(a^i; e^{*,i}), \text{ a.e. } (a^i, l^i, \epsilon^i) \in \text{Support}(\lambda^{*,i}), \tag{10}$$

where $e^{*,i} = (l^i, \epsilon^i, w^{\theta^*}(k^*), R^{\theta^*}(k^*), T^{*,i})$ and $e^{**,i} = (l^i, \epsilon^i, w^{\theta^{**}}(k^*), R^{\theta^{**}}(k^*), T^{**,i})$. If the inequality is reversed, then household i's direct response is instead *negative*.

A couple of comments on notation are useful here. First, we condition the saving correspondences on the environment θ to emphasize its potential shifts in response to changes in this environment. Second, notice that in $e^{*,i}$ and $e^{**,i}$, factor prices and transfers are allowed to change because the environment changes, but are evaluated at the same capital-labour ratio, k^*, highlighting the partial equilibrium nature of the exercise here—hence the emphasis on "direct." Finally, if the saving correspondence is not single-valued, then the inequality in (10) refers to the strong set order, that is, the least and greatest optimal savings levels must increase. This convention is adopted throughout the rest of the paper.

Definition 5 (Direct responses)**.** Let k^* be an equilibrium given the environment $\theta^* \in \Theta$ and consider a different environment $\theta^{**} \in \Theta$. We say that the *direct response is positive* if the mean asset holdings of households increase at k^* when the environment changes from θ^* to θ^{**}, *i.e.* if $\int \hat{a}^{\theta^{**},i}(k^*)\, di \geq \int \hat{a}^{\theta^*,i}(k^*)\, di$. If the inequality is reversed so that the mean asset holdings decrease at k^* when the environment changes from θ^* to θ^{**}, the *direct response is negative*.

The definition is intuitive: We average over the asset holdings (or gross savings) of households in the old and new environments holding the capital-labour ratio k^* (hence prices) fixed, and trace the direction of change. As we illustrate in Section 4, the definition makes direct reference to the associated consumption-savings model. In particular, for given k^*, the relevant asset holdings can be computed without any knowledge of (general) equilibrium changes in prices or quantities that follow from the change in environment. Clearly, if individual direct responses in Definition 4 are uniformly positive, the (aggregate) direct response in Definition 5 is positive.

Note that in both Definitions 4 and 5, (pre-tax) *market prices* are fixed at their initial steady-state values. For example, if the only change in environment is a change in the capital tax rate ($\theta = \tau$), then we have $w^{\theta^{**}}(k^*) = w^{\theta^*}(k^*) = f(k^*) - f'(k^*)k^*$, and $R^{\theta^*}(k^*) = (1-\tau^*)f'(k^*) - \Delta$ and $R^{\theta^{**}}(k^*) = (1-\tau^{**})f'(k^*) - \Delta$. So when investigating whether a change in environment leads to a positive or negative direct response, it is sufficient to consider the consumption-savings problem in steady state, with given prices. By comparison, the standard approach in the robust comparative statics literature—including in our own work, Acemoglu and Jensen (2015)—is to impose positive direct responses in the sense of Definition 4 uniformly across all households and for all market prices (all capital-labour ratios).[17] In Section 4, we illustrate how the direction of the direct response can be determined in growth models with quasi-hyperbolic preferences, self-control and temptation utilities and systematic misperceptions, and in all of these cases such results are made possible by the fact that we only need to determine the direction of the direct response, without taking into account any general equilibrium changes in prices.

We can now state the simplest version of our main result, which establishes that the long-run equilibrium outcome is pinned down by the direct response.

Theorem 1 (Main theorem, unique steady state)**.** *Assume that households satisfy the assumptions in Lemma 1. For environments $\theta^*, \theta^{**} \in \Theta$ let k^* and k^{**} denote associated non-trivial steady-state equilibria and assume that these are unique. Then $k^{**} \geq k^*$ if and only if the direct response is positive when the environment changes from θ^* to θ^{**}. Similarly, $k^{**} \leq k^*$ if and only if the direct response is negative when the environment changes from θ^* to θ^{**}.*

Although uniqueness is a special case, the theorem captures this paper's main message: In one-sector growth models, long-run outcomes are entirely pinned down by the average of the *direct responses*. Misperceptions, biases, and other departures from standard, fully rational, and time-separable preferences thus impact long-run outcomes in so far as they influence household decisions at given prices. This result also implies that such departures can easily lead to "paradoxical" comparative statics (which reverse those of the standard neoclassical growth model) provided that they change the sign of the direct response. Conversely, when they do not do so, despite the very rich and potentially complex general equilibrium interactions that these behavioural preferences may spawn, they will not affect the qualitative properties of the long-run equilibrium. In the next section, we use this theorem in economies with quasi-hyperbolic preferences, self-control and temptation utilities and systematic misperceptions to investigate the direction of comparative statics (how our results can be applied with other classes of behavioural preferences and biases is discussed in online Appendix C (supplementary material)).

Theremainder of this subsection generalizes Theorem 1 to situations with multiple equilibria and extends the discussion of the intuition and the mathematical arguments from the introductory section. We next show that both necessity and sufficiency in our main result remain valid when there are multiple equilibria provided that we focus on the least or the greatest steady state and the exogenous changes we are considering are "small" (meaning that we can choose them to be small enough in the usual implicit function theorem sense).

Theorem 2 (Greatest and least steady states under multiplicity I)**.** *Assume that households satisfy the assumptions in Lemma 1 and let $k^*_- = \inf\{k : k \in \mathcal{M}^{\theta^*}(k)\}$ denote the least steady state and $k^*_+ = \sup\{k : k \in \mathcal{M}^{\theta^*}(k)\}$ the greatest steady state when the environment is $\theta^* \in \Theta$, and analogously k^{**}_- and k^{**}_+ when the environment is $\theta^{**} \in \Theta$. Assume in addition that $\mathcal{M}^\theta$ is upper hemi-continuous in $\theta \in \Theta$ (where now Θ is a topological space). Consider an infinitesimal change in the environment to θ^{**}. Then, $k^{**}_- \geq k^*_-$ if and only if the direct response is positive at k^*_- when the environment changes from θ^* to θ^{**}, and $k^{**}_+ \geq k^*_+$ if and only if the direct response is positive at k^*_+ when the environment changes from θ^* to θ^{**}.*

17. For example, in the Bewley–Aiyagari model, one can use the results in Light (2020) who establishes that households will increase their savings if preferences are CRRA, the coefficient of relative risk aversion is less than one, and the rate of return increases (see his Theorem 1). In contrast, we will not impose such uniform positive or negative direct responses. Rather, our approach relies on just the sign of the direct response at the (initial) steady-state capital-labour ratio k^*, the direct response is positive (or negative).

If there are multiple equilibria and the change in environment is not "small" (or we are unwilling or unable to place a topology on the set of possible environments Θ), the sufficiency part of our main result will still hold for the greatest equilibrium when the direct response is positive (and for the least equilibrium when the direct response is negative):

Theorem 3 (Greatest and least steady state under multiplicity II)**.** *Assume that households satisfy the assumptions in Lemma 1 and consider* $k^* = \sup\{k : k \in \mathcal{M}^{\theta^*}(k)\}$ *(the greatest steady state) of the environment* $\theta^* \in \Theta$*. Then if the direct response is positive at* k^* *when the environment changes from* θ^* *to a new environment* $\theta^{**} \in \Theta$*, the economy's greatest steady state increases, i.e.* $\sup\{k : k \in \mathcal{M}^{\theta^{**}}(k)\} \geq k^*$*. Analogously, consider* $k^* = \inf\{k : k \in \mathcal{M}^{\theta^*}(k)\}$ *(the least steady state) of the environment* $\theta^* \in \Theta$*. Then if the direct response is negative at* k^* *when the environment changes from* θ^* *to the new environment* $\theta^{**} \in \Theta$*, the economy's least steady state decreases, i.e.* $\inf\{k : k \in \mathcal{M}^{\theta^{**}}(k)\} \leq k^*$*.*

Appendix B contains additional results along the lines of the previous two theorems. Although important for theoretical applications, the details are less central to our substantive results, hence their relegation to the Appendix. In addition, we also provide there a detailed comparison with the related equilibrium comparative statics results in Milgrom and Roberts (1994) and Acemoglu and Jensen (2013).

The intuition for the results presented in this section was already discussed in the Introduction. Here, we had elaborate their mathematical and conceptual underpinnings. Most importantly, our approach enables us to represent any model that falls within the general one-sector behavioural growth model with a *market correspondence* $\mathcal{M}^{\theta}$. From Lemmas 1 and 2, saving correspondences have no jumps down which, which implies that the market correspondence will be compact- and convex-valued, upper hemi-continuous and begin above and end below the 45° line. Crucially, a positive direct response will raise (or "shift up") the market correspondence at the initial capital-labour ratio k^* as illustrated in Figure 2 (this is proved in the

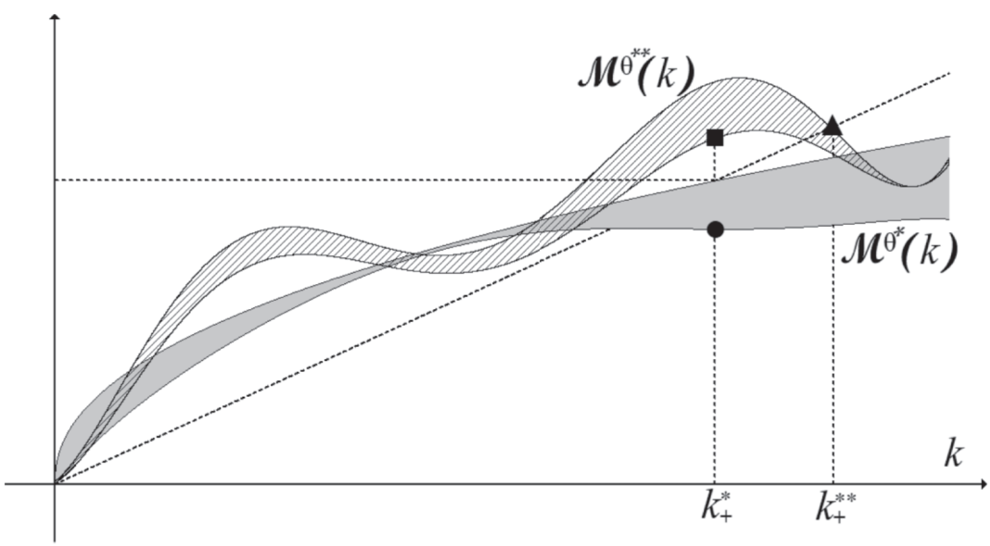

FIGURE 2

A positive direct response shifts the market correspondence up at k_+^* (shown by the move from the dot to the box) and leads to a higher steady-state capital-labour ratio (shown by the triangle). The Figure depicts a case in which there are multiple steady states both before and after the change in environment key Lemma 4 in the Appendix). As this figure also illustrates, the new equilibrium k^{**} must then be above k^*, regardless of whether the market correspondence shifts up or down at other capital-labour ratios $k \neq k^*$. This result implies that the direct response of aggregate savings at k^* pins down the direction of change for the steady state.[18]

4. APPLICATIONS

In this section, we provide a number of applications of our general framework. Throughout our emphasis will be on two aspects. First, we show that applying our results is often quite straightforward. Second, we establish that even simple applications of these methods lead to a number of new results relative to the existing literature. To ease interpretation, we work with models that have a lot of commonality. Specifically, in all cases, we start from Bewley–Aiyagari–style models with incomplete markets, in which households receive shocks to their labour income or endowments.

The next subsection presents a lemma that characterizes how solutions of (generalized) Euler equations change in response to variation in environment. This result will be used in some of our applications and is of independent interest. We then provide theoretical results for three classes of models. The first are those that contain "systematic misperceptions" about future variables, such as interest rates or labour income. We show that our methods can be applied readily in this class of models. Second, we turn to models of self-control and temptation utility, as developed in Gul and Pesendorfer (2004), and explain how our results lead to a number of new results in this context. Third, we discuss models of quasi-hyperbolic discounting as in Phelps and Pollak (1968), Laibson (1997), Harris and Laibson (2001), and Laibson *et al.* (2020). In the last part of this section, we show how our results can be blended with numerical methods in order to obtain additional insights. Throughout this section, the emphasis is on the direct response of an economy to changes in environment—how equilibrium objects change holding constant market prices. With such a characterization at hand, our main theorems can be invoked to establish general comparative static results.

4.1. *A useful lemma: shifts of solutions to Euler equations*

Our main results in this paper rely on characterizing direct responses in the sense of Definition 5. In deterministic environments with appropriate smoothness and boundary conditions, these direct responses can be obtained from (steady state) Euler equations. In stochastic environments, there is typically no such simple Euler equation. Nevertheless, our main result in this subsection, Lemma 3, shows how various changes in the environment shift the set of solutions to stochastic (and potentially generalized) Euler equations. These shifts can then be combined with Theorems 1–3 to derive equilibrium comparative statics in some of the applications we consider.

To illustrate our approach, let us start with the benchmark Bewley–Aiyagari model and, as we will do throughout this section, let us suppress dependence on factor prices and transfers, assume differentiability, and write the time-stationarity saving function as $s(a, l, \epsilon)$. In the Bewley–Aiyagari model, households are only uncertain about their future labour endowments so there is no loss of generality in omitting ϵ and writing the time-stationary savings function $s(a, l)$. Then

18. The figure illustrates the "general" case, in the sense that there are multiple steady states both before and after the change in the environment, and we focus on the largest ones corresponding to k^*_+ and k^{**}_+ in Theorems 2 and 3.

the next period's assets choice $a' = s(a, l)$ solves a "Deaton-type" Euler equation:[19] for *a.e.* $(a, l) \in [\underline{a}, \overline{a}] \times [l_{\min}, l_{\max}]$,

$$\begin{aligned} &L(a', (a, l), s, \rho) \\ &\quad = -u'((1+R)a + wl - a') \\ &\qquad + \max\{\delta(1+R) \int u'((1+R)a' + wl' - s(a', l'))\mu(dl'), u'((1+R)a + wl - \underline{a})\} = 0 \end{aligned}$$

Here, ρ summarizes all of the fixed parameters on the right-hand side (the prices R and w, and the environment $\theta \in \Theta$ including the borrowing constraint $\underline{a}$). As we explain in the rest of this section, stochastic Euler equations in several other behavioural consumption-savings models can be written in a similar form. In this spirit, let us define a time-stationary saving function for household $i \in [0, 1]$ in the general behavioural growth model as a solution $a'^{,i} = s^i(a^i, l^i, \epsilon^i)$ to the (steady state) Euler equation:

$$L^i(a'^{,i}, (a^i, l^i, \epsilon^i), s^i, \rho^i) = 0 \quad \text{for } a.e. (a^i, l^i, \epsilon^i) \in [\underline{a}^i, \overline{a}^i] \times [l^i_{\min}, l^i_{\max}] \times E^i. \tag{11}$$

It is clear that the Bewley–Aiyagari model is a special case. In general, (11) can easily have multiple solutions, and if so, we say that a solution is the least (resp., greatest) solution if the level of savings is weakly below (resp., weakly above) the level of savings of any other solution to (11) for all $(a^i, l^i, \epsilon^i) \in [\underline{a}^i, \overline{a}^i] \times [l^i_{\min}, l^i_{\max}] \times E^i$

From now on, we fix a specific household $i \in [0, 1]$ and omit the index i.

Assumption 2.

(1) ***Continuity***: *$L(a', (a, l, \epsilon), s, \rho)$ is continuous in $a' \in [\underline{a}^i, \overline{a}^i]$ and $(a, l, \epsilon) \in [\underline{a}, \overline{a}] \times [l_{\min}, l_{\max}] \times E$.*

(2) ***Boundary conditions:*** *Given s and ρ, $L(\underline{a}, (a, l, \epsilon), s, \rho) \geq 0$ and $L(\overline{a}, (a, l, \epsilon), s, \rho) < 0$ for all $(a, l, \epsilon) \in [\underline{a}, \overline{a}] \times [l_{\min}, l_{\max}] \times E$.*

(3) ***Monotonicity in future savings***: *Given ρ, $L(a', (a, l, \epsilon), s, \rho) \leq L(a', (a, l, \epsilon), \tilde{s}, \rho)$ if $\tilde{s}(a, l, \epsilon) \leq s(a, l, \epsilon)$ for all $(a, l, \epsilon) \in [\underline{a}, \overline{a}] \times [l_{\min}, l_{\max}] \times E$.*

The first two parts of the assumption impose weak regularity conditions and are satisfied in all of our applications. In particular, the boundary conditions hold in all of our applications under an (upper) Inada condition on utility.[20] The third part is the key monotonicity condition, which is a restriction on underlying parameters and functional forms. This third part is also satisfied in the benchmark Bewley–Aiyagari model and holds in our applications, except in the "sophisticated" quasi-hyperbolic model, where we will not use this approach (see Section 4.4).

The next lemma shows how least and greatest solutions change as we modify the environment.

Lemma 3. *Suppose that Assumption 2 holds. Then (11) has least and greatest solutions. Furthermore:*

- *If s^* is the unique solution to (11) when $\rho = \rho^*$, s^{**} is the unique solution to (11) when $\rho = \rho^{**}$, and $L(s^*(a, l, \epsilon), (a, l, \epsilon), s^*, \rho^{**}) \geq L(s^*(a, l, \epsilon), (a, l, \epsilon), s^*, \rho^*)$ for $a.e.$ (a, l, ϵ) in the support of λ_{s^*}, then $s^{**}(a, l, \epsilon) \geq s^*(a, l, \epsilon)$ for all (a, l, ϵ).*

19. See Deaton (1991) and Li and Stachurski (2014).

20. In the benchmark Bewley–Aiyagari model, if $u'(c) \to 0$ as $c \to \infty$, we can always find an upper bound on assets such that if a' is above this bound then $L < 0$. Moreover, we have $L \geq 0$ if $a' = \underline{a}$, since otherwise $u'(Ra + wl - \underline{a}) > \delta R \int u'(R\underline{a} + wl' - s(\underline{a}, l'))\mu(dl') > u'(Ra + wl - \underline{a})$, which is impossible.

- *The statement remains valid for the least and greatest solutions, even when there are multiple solutions, provided that the change from ρ^* to ρ^{**} is infinitesimal and L is continuous.*
- *The statement also remains valid for the greatest solutions s^* and s^{**}, even when there are multiple solutions.*

This lemma shows that we can determine whether a change in the environment leads to a positive direct response without imposing the usual conditions for monotone comparative statics, and also without having to compute the saving function in the new environment, which can be challenging in many economies with shocks and behavioural biases. Crucially, the key condition in Lemma 3 only needs to hold given ρ^* and s^*—that is, at the initial solution to the generalized Euler equation, s^*, given the initial parameters and steady-state prices w and R and the initial invariant distribution λ_{s^*}. The lemma thus allows for local (partial equilibrium) analysis, as required for our main results.

As a final remark, we note that the lemma is stated analogously to Theorems 1–3, distinguishing cases with a unique solution from those with multiple solution with or without small shocks. In principle, all of the results in this section should be stated in this manner. However, with an abuse of mathematical precision, in what follows, we simplify the statements of our results by writing simply "a change from ρ^* to ρ^{**} increases the saving function." Throughout, this should be interpreted as either applying under conditions of uniqueness or for the least and the greatest steady-state equilibria under the appropriate conditions as in Theorems 1–3 or Lemma 3.

4.2. *Systematic misperceptions*

Our first application is to economies with systematic misperceptions, where agents may not use the "true model" or may make other systematic mistakes in forming their expectations. In our general formulation, household decisions depend on beliefs about both future (exogenous) variables and about future prices, which then shape expectations about future selves' savings. In this subsection, we allow for misperceptions on all three dimensions: (i) agents may persistently overestimate their future level of patience (the discount factor), in which case they will systematically overestimate their future savings; (ii) they may persistently overestimate their future labour income (which will directly impact savings decisions today); and (iii) they may, alternately, believe that some policies, such as changes in the capital income tax rate, will affect their labour income in ways that are not consistent with the underlying model. In all of these cases, there is a natural dynamic inconsistency: consumption and saving plans made with incorrect beliefs will have to be revised once households are confronted with actual realizations.

As noted above, we consider as benchmark a Bewley–Aiyagari model with ex-ante identical households, subject to *i.i.d.* labour endowment shocks, given by $l_t \sim \mu(\cdot)$ over a bounded support $[l_{\min}, l_{\max}] \subseteq \mathbb{R}_{++}$. We maintain the same assumptions on borrowing limits as in the previous subsection. All households have time-separable and neoclassical (continuous, increasing, strictly concave, satisfying Inada conditions) utility given by u, and geometrically discount the future with discount factor $\delta < 1$. We assume that a fraction $1-\alpha \in [0,1]$ of households are "rational": they know the true distribution of l_t and form correct beliefs about all future variables. The remaining fraction α are "behavioural" and may hold persistently wrong beliefs.

The first set of incorrect beliefs we consider relate to future patience/discount factors. Specifically, while the true discount factor is always δ, behavioural households overestimate (underestimate) their future patience if they believe that future selves will base their decision on the discount factor $\hat{\delta} > \delta$ ($\hat{\delta} < \delta$).[21] These beliefs are assumed to be time-invariant and

21. Formally, the parameter ϵ^i introduced in Section 2.2 now parameterizes beliefs about future discount factors: $\epsilon^i = \delta^i$.

hence dogmatic, in the sense that a behavioural individual does not change her beliefs even after realizing that her expectations so far have not been realized.[22] Denoting individual i's beliefs about the marginal distribution of her future selves' discount factor δ^i by $\hat{P}_t^i$, we write $\hat{P}_t^i(\delta^i = \hat{\delta}|e_t^i) = 1$ for all t, for all $i \in [0, 1]$ and for all e_t^i, where e_t^i summarizes factor prices, policies and endowments (see the proof in the Appendix for details).

Proposition 1. *Suppose that a fraction $\alpha \in [0, 1]$ of households systematically overestimate (resp., underestimate) their future patience. Then, the steady-state capital-labour ratio increases (decreases) when α increases.*

Several points are worth emphasizing. First, as noted above, with an abuse of mathematical precision, the statement refers to the steady-state capital-labour ratio increasing or decreasing. This should be read as either applying under uniqueness,[23] or applying for small changes (in arbitrarily small increase in α), or as referring to the greatest steady-state capital-labour ratio.

Second, this proposition heavily relies on Lemma 3. The proof first applies this lemma to show that behavioural households that overestimate future patience save more than fully rational households. Once this result is established, Theorems 1–3 yield the desired conclusions. This structure of argument also shows that our methods are in fact quite straightforward to apply in this class of environments.

Third, this proposition also determines the effects of misperceptions relative to the neoclassical benchmark: with households that systematically overestimate their patience, steady-state capital-labour ratio is higher than in the fully rational benchmark (which corresponds to the special case where $\alpha = 0$).

Fourth, this result is, at some level, intuitive. When a household overestimates their future patience, they think they will have higher savings and thus lower consumption in the next period. This implies, from the concavity of u, that they will overestimate the marginal utility of future consumption, encouraging them to save more. While there are indirect effects on their saving behaviour, for example, coming from the implications of the life-time stochastic budget constraint, this marginal utility channel is strong enough to ensure that an economy with more behavioural households has more savings and thus a higher steady-state capital-labour ratio.

Fifth, although it follows from a direct application of our methods, to the best of our knowledge there are no analogues of this type of result in the literature. In fact, there are several competing equilibrium effects, which make the impact of such a change on the steady state quite complex. To understand this point, note that an increase in α raises savings, as described in the previous paragraph, and consequently reduces the interest rate and increases wages. These price changes will have ambiguous implications for both the behavioural and rational households. Depending on the income and substitution effects, the equilibrium response to price changes may be a further increase or a reduction in savings. As a result, there is no reason to expect that the full general equilibrium effect will go in the same direction as the direct effect. The finding that it does so under general conditions is an original result of our framework, which highlights the critical role of the one-sector structure (see the discussion in the Introduction as well).

Sixth, the proposition says nothing about individual-level behaviour. In fact, we will see at the end of this subsection that there is generally a type of individual-level "indeterminacy" (see

22. Such dogmatic beliefs are important, since otherwise Bayesian updating would lead to changes in beliefs after a sufficiently long sequence of realized labour incomes. For a discussion of how these types of beliefs may survive long sequences of contradictory information, see Benjamin *et al.* (2015).

23. In this case, it is possible to place stronger conditions to guarantee uniqueness. For example, see Light (2020) and Light and Weintraub (2021) (Section 5), whose conditions are sufficient to ensure uniqueness in our model, despite the systemic misperceptions of some households.

Proposition 4): because of equilibrium responses to changes in prices, some households will end up increasing their savings while others will reduce theirs, and it is very difficult to pin down how a given household will behave, without knowing the exact changes in equilibrium prices (which of course depends on how each household behaves at the end). This indeterminacy and the resulting richness of individual behaviour sharply distinguish our approach from those that use monotone comparative static methods that require all households to move in the same direction.

Finally, this proposition, like all others in this section, is stated for the case of *ex ante* homogeneity in terms of utility functions and labour endowment sequences (but of course not in terms of rationality/behavioural biases). This homogeneity is adopted for simplicity and can be easily relaxed. We could allow, as in our main analysis, different types of households, with each group having different utility functions and different fractions of behavioural and fully rational agents. In that case, an analogue of Proposition 1 follows, provided that we consider a change in the environment that still induces a positive direct response as required in Definition 5. In fact, using similar steps, one can also combine different types of behavioural biases within the same model, and if the change in environment induces a positive direct response, our main results can be readily applied.

The results in Proposition 1 critically depend on the fact that we are considering misperceptions about future discounting. If, instead, there are misperceptions about labour income, the results are very different. To illustrate this possibility, let us now suppose that a fraction α of the households believe that their labour income has a distribution given by $w_{t'}l_{t'} \sim \mu_W^m(\cdot|w_{t'})$ where μ_W^m may differ from the true distribution of labour income, $\mu_W(A|w_{t'}) = \mu\{l \in [l_{\min}, l_{\max}] : w_{t'}l \in A\}$ (here A is a Borel measurable subset of $\mathbb{R}_+$). In what follows, we say that behavioural households "overestimate (resp., underestimate) their future labour income" given the market wage w, if, given w, $\mu_W^m(\cdot|w)$ first-order stochastically dominates (resp., is first-order stochastically dominated by) $\mu_W(\cdot|w)$. Once again, these beliefs are dogmatic.

Proposition 2. *Suppose that a fraction $\alpha \in [0, 1]$ of households systematically overestimate (resp., underestimate) their future labour income at the initial steady-state wage level $w(k^*)$ where k^* is initial steady-state capital-labour ratio. Then, the steady-state capital-labour ratio decreases (increases) when α increases.*

To save space, the proofs of this and the remaining results in the paper are presented in the online Appendix C (supplementary material).

Although households are again overestimating future savings, the conclusions are the opposite of Proposition 1: behavioural biases now reduce savings and capital accumulation. This is because incorrect beliefs about labour income have very different implications than those about future discount rates, as they encourage households to consume more under the mistaken belief that they are richer than they truly are. As a result, greater overestimation of future labour income leads to lower savings, and the behavioural model has, analogously, lower capital-labour ratio than the fully rational benchmark (which again corresponds to the case where $\alpha = 0$).

We would like to reiterate that, despite the apparent simplicity of this result, we are not aware of similar findings in the literature. In fact, an approach that focuses on aggregate behaviour is key for deriving this result, since typically some individuals will increase their savings while others reduce theirs.

Our next result shows that simple misperceptions can also change the direction of standard neoclassical comparative statics. We illustrate this possibility focusing on one of the more robust comparative statics in fully rational models: the positive impact of a reduction in the capital income tax rate on capital accumulation.

Suppose now that capital income is taxed at the rate $\tau \in [0, 1)$, there is no tax on labour income, capital depreciates fully after use, and tax revenues are spent on a non-productive public good (and thus do not impact the marginal utility of consumption). In terms of our general formulation, this implies $G(k_t) = \tau f'(k_t)k_t$ and $T^i = 0$ for all i, and equilibrium prices are given by (1)–(2) where $\omega(k_t) = 0$, $\tau(k_t) = \tau$ and $\Delta = 1$. For illustration purposes, let us focus on CRRA utility and further assume that the intertemporal elasticity of substitution χ is greater than some threshold $\underline{\chi} \in (0, 1)$, which ensures that the substitution effect is not overwhelmed by the income effect and thus households' asset supply is increasing in the rate of interest (see Aiyagari, 1994, pp. 667–668). These assumptions are sufficient to ensure that a reduction in the capital income tax rate increases the (unique) steady-state capital-labour ratio in the benchmark Bewley–Aiyagari model.

The only difference between the rational and behavioural households is that the behavioural households incorrectly believe that capital income taxes directly impact their future labour incomes (rather than just indirectly via capital accumulation). In other words, they believe: $w_{t'}l_{t'} \sim \mu_W^m(\cdot|w_{t'}, \tau)$, $t' > t$. We say that a reduction in capital income taxes "causes optimism (resp., pessimism)" if $\mu_W^m(\cdot|w_{t'}, \tilde{\tau})$ first-order stochastically dominates (resp., is first-order stochastically dominated by) $\mu_W^m(\cdot|w_{t'}, \tau)$ whenever $\tilde{\tau} < \tau$. For example, a reduction in capital income taxes can cause optimism if some agents believe that such a reduction increases the efficiency of the economy beyond its impact on saving incentives.

Proposition 3. *Suppose that all households have CRRA utility with intertemporal elasticity of substitution $\chi \in (\underline{\chi}, \infty)$ and a fraction $\alpha \in (0, 1]$ of households systematically misperceive the effect of capital income taxes on their future labour income. If a reduction in capital income tax causes pessimism among behavioural households, then it increases the steady-state capital-labour ratio. If it instead causes optimism among behavioural households, then for any α there exists $\chi_a \in (\underline{\chi}, \infty)$ such that for all $\chi \leq \chi_\alpha$, the lower capital income tax reduces the steady-state capital labour ratio.*

The proof of this result relies on Lemma 3 as well as Theorem 1 in Light (2020).

When the capital income tax is reduced, rational households always (for any choice of $\chi \in (\underline{\chi}, \infty)$) increase their savings starting from the initial steady state. If the capital income tax reduction causes pessimism among behavioural households, their reaction will amplify the response relative to the benchmark with rational households (this is because behavioural households feel poorer and thus increase their savings by even more than the rational agents). The aggregate direct response of Definition 5 is positive, and the standard comparative statics hold by Theorem 1 (we can invoke this theorem because we have uniqueness in this case).

However, if the capital income tax reduction causes optimism among behavioural households, their response at the initial prices can be negative. In fact, when χ is sufficiently low, because their response to the change in the net interest rate is very small, the optimism channel wins out and they respond negatively to the cut in the capital income tax rate. Hence, in this case the direct responses of rational and behavioural households are going in opposite directions, and the balance between the two will depend on their quantitative magnitudes. When χ is low, rational households' response is quantitatively small, and thus the negative reaction from behavioural households wins out and we obtain the second part of the proposition, which shows the reversal of the neoclassical comparative statics in response to capital tax rates. Notably, even a small fraction of behavioural agents that mistakenly become more optimistic about their future labour income is sufficient for such a reversal.

This proposition illustrates how fine details of behavioural biases are necessary to understand whether standard comparative statics will continue to apply. It also reiterates why monotonicity-based tools would not have been useful in the setup (different households are moving in different directions).

Our final result takes this one step further and establishes individual-level indeterminacy, as already anticipated above. Let $\eta(\cdot)$ denote the Lebesgue measure on the set of households [0, 1] so that $\eta(J)$ is the mass of a (measurable) subset of households $J \subseteq [0, 1]$.

Proposition 4. *Suppose that each household has CRRA utility with intertemporal elasticity of substitution $\chi_i \in (1, \infty)$ and consider a reduction in the capital income tax that causes pessimism among the behavioural subset of households. Then there exists $B > 0$ such that the following holds: For any (measurable) subset $J \subseteq [0, 1]$ of households with $\eta(J) \leq B$, there exists a production function and a profile of misperceptions and intertemporal elasticities of substitution for the remaining set of households, $[0, 1]\backslash J$, such that the lower capital income tax will lead to lower aggregate stationary savings for all households in J, while the steady-state capital-labour ratio and aggregate savings will increase.*

For the same production function, there also exists a profile of misperceptions and intertemporal elasticities of substitution for the remaining set of households, $[0, 1]\backslash J$, such that the lower capital income tax leads to higher aggregate stationary savings for all households in J.

Focusing on the comparative statics with respect to the capital tax rate, Proposition 4 shows that, while aggregate savings increase, there is not much that can be said about individual behaviour. In particular, any small subset J of households will increase or reduce their savings depending on the exact misperceptions and elasticities of substitutions of other agents.[24] The intuition for this result is that for any level of the interest rate elasticity of households in the subset J, the remaining households' saving levels could be even more elastic. This would make the increase in the after-tax interest rate small relative to the rise in the wage rate, and the resulting large income effect induces households in J to reduce their savings. This indeterminacy result reiterates that our main results are not driven by some hidden monotonicity assumptions—they are a consequence of the discipline that this class of models imposes on aggregate variables despite, despite behavioural preferences, while placing little or no restrictions on individual behaviour.

4.3. *Self-control and temptation*

We next present similar comparative static results for self-control and temptation preferences introduced and studied in Gul and Pesendorfer (2004). The benchmark is as before; a Bewley–Aiyagari model with *ex ante* identical households subject to *i.i.d.* labour endowment shocks, given by $l_t \sim \mu(\cdot)$ over support $[l_{\min}, l_{\max}] \subseteq \mathbb{R}_{++}$. The main difference is that now, in addition to a standard neoclassical utility function u and discount factor $\delta < 1$, households have a temptation cost given by ϕv where the parameter $\phi \in [0, 1]$ represents their "temptation intensity." As $\phi \to 0$, we approach the standard neoclassical benchmark without self-control and temptation problems. We continue to assume a borrowing limit of $\underline{a} \leq 0$, and to start with, there are no misperceptions. In this case, we assume, again for simplicity, that all households have self-control and temptation preferences, rather than doing so only for a fraction α of households. We also suppress ϵ in what follows to pare down the notation.

This model satisfies the assumptions in Lemma 1, provided that overall utility, $u(c) + \phi v(c)$, is concave, increasing and continuous. To simplify the exposition, we will additionally assume that u is strictly concave and that u and v are at least four times continuously differentiable

24. An analogous result holds when aggregate savings decrease following the decline of the capital income tax rate. Here, for simplicity, we focus on the more standard case in which aggregate savings and the capital-labour ratio increase.

on $\mathbb{R}_+$. In addition, we assume that v is either strictly convex everywhere, or strictly concave everywhere with positive third derivative.

The advantage of this formulation of self-control and temptation, as introduced in Gul and Pesendorfer (2004), is that when prices are constant, then consumption-saving decisions are given by a standard dynamic programming problem. In particular, households' time-stationary saving function $s(a,l)$ is uniquely determined, and in the rest of the section, we simplify the notation further by dropping the conditioning on factor prices. In this case, the following dynamic programming recursion determines the saving function $s(a,l)$:

$$s(a,l) = \arg \max_{\underline{a} \leq a' \leq y} u(y-a') + \phi v(y-a') + \delta \int W((1+R)a' + wl')\mu(dl') - \phi v(y - \underline{a}) \,, \tag{12}$$

where $y = (1+R)a_t + wl_t$ denotes current total wealth (or cash-at-hand) and W is the value function.

Our first result shows the effects of changing the temptation intensity ϕ or the borrowing limit $\underline{a}$.

Proposition 5. *Suppose that the steady-state saving function in this economy $s(a,l)$, and assume that*

$$\delta(1+R) \int \frac{v'((1+R)s(a,l) + wl' - s(s(a,l)), l') - v'((1+R)s(a,l) + wl' - \underline{a})}{v'((1+R)a + wl - s(a,l))} \mu(dl') \leq 1. \tag{13}$$

Then the steady-state capital-labour ratio is decreasing in the temptation intensity ϕ. If, on the other hand, this inequality is reversed, then the steady-state capital-labour ratio is increasing in the temptation intensity ϕ.

Suppose $\phi > 0$. Then, a looser borrowing constraint (a reduction in $\underline{a}$) will reduce the steady-state capital-labour ratio if v is convex; and it will increase the steady-state capital-labour ratio if v is concave and no household is initially borrowing constrained.

Once again, we are not aware of any results in the literature that are similar to this proposition, which illustrates how our general approach can be applied to yield simple but powerful new insights. Although far from obvious, these results are intuitive. Condition (13) implies that costly self-control does not raise the (expected) marginal utility of future consumption by "too much" in comparison with the benchmark case. In particular, this is true in the case (Gul and Pesendorfer, 2004) focus on, since a convex temptation cost function v ensures that self-control *reduces* the marginal utility of future consumption when ϕ increases. As this will induce households to shift consumption towards the present, savings decline as ϕ increases. Our main theorems then imply that a higher ϕ leads to a greater steady-state capital-labour ratio, as illustrated in Figure 3. If, on the other hand, costly self-control increases the marginal utility of future consumption (v concave), and this impact is sufficiently powerful (v is "sufficiently concave"), then the impact of the higher ϕ on savings and, by our main theorems, the new steady state is reversed.[25] It then follows immediately that self-control and temptation preferences can increase or reduce

25. We can also note that convex temptation costs reduce the "over-saving" problem in the benchmark Bewley–Aiyagari model, while concave temptation costs exacerbate it.

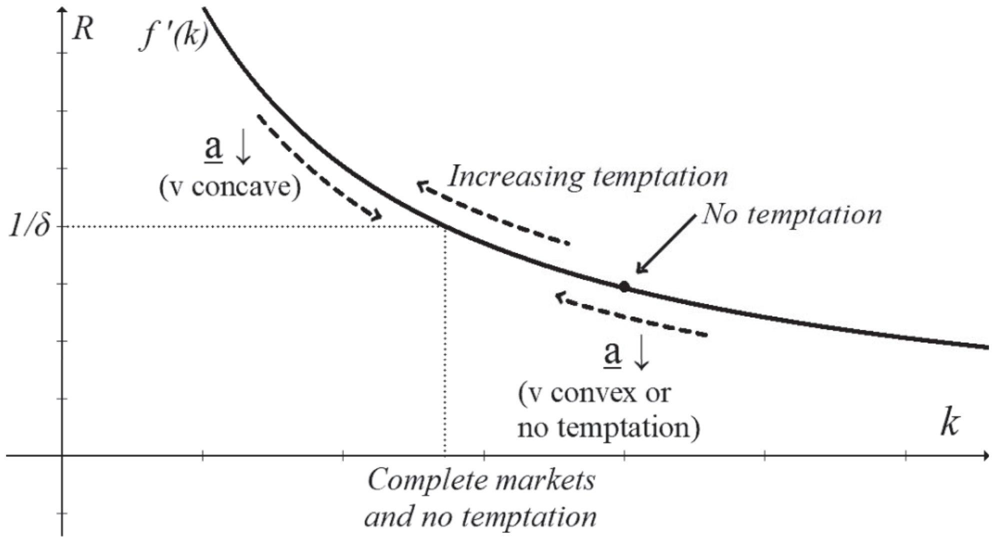

FIGURE 3
Temptation versus Benchmark Case

steady-state capital-labour ratios relative to the neoclassical benchmark (with $\phi = 0$), depending on whether condition (13) is satisfied.

The second part of the proposition might at first appear even more surprising. Recall that temptation costs reduce the marginal utility of future consumption when v is convex because the marginal cost of giving in to temptation increases with the household's wealth level. With a similar reasoning, looser borrowing constraints that allow the household to go into debt increase temptation costs when v is convex. This reduces the marginal utility of future consumption and encourages lower savings. In this case, the comparative statics are thus similar to the standard ones (Aiyagari, 1994, p. 672), but working through a distinct temptation channel. In contrast, when v is concave, households will become better at resisting temptation with a looser borrowing constraint, because temptation increases less than proportionately with household wealth.[26] In this case, paradoxically, a looser borrowing constraint can increase aggregate savings, in particular for households that are not actually borrowing constrained.[27]

We next study the effects of tax policy in the presence of self-control and temptation. We now set $\phi = 1$ to economize on notation. The only additional feature is that, as in the previous subsection, there is a linear capital income tax rate at the rate τ, the proceeds of which are spent on a non-productive public good. We also assume full depreciation, so that $1 + R = (1 - \tau)\hat{R}$ where $\hat{R} = f'(k^*)$ is the (pre-tax) market price. To shorten expressions, we further use the notation $[u' + v'](c) = u'(c) + v'(c)$ and similarly the second derivatives.

26. Loosely speaking, we can think of this as the case in which the marginal temptation to eat a (whole) cake is getting weaker with the size of the cake due to diminishing returns.

27. Going against this is the fact that households that are borrowing constrained will increase their consumption (see *e.g.* Aiyagari, 1994, p. 672). For this reason, this part of the proposition is stated for the case in which no household is initially borrowing constrained.

Proposition 6. *Denote the initial rate of capital income tax by $\tau^* \in (0,1)$, let $s(a,l)$ denote the initial saving function and $c(a,l) = (1+R)a + wl - s(a,l)$ the consumption function. If*

$$\delta(1-\tau^*)\hat{R}s(a,l)\frac{\int [u''+v''](c(s(a,l),l') - v''(\theta(1-\tau^*)\hat{R}s(a,l)+wl'-\underline{a}))\mu(dl')}{\int [u'+v'](c(s(a,l),l') - v'((1-\tau^*)\hat{R}s(a,l)+wl'-\underline{a})\mu(dl')}$$
$$\geq \frac{[u''+v''](c((a,l))a}{\int [u'+v'](c(s(a,l),l') - v'((1-\tau^*)\hat{R}s(a,l)+wl'-\underline{a})\mu(dl')} - \delta \qquad (14)$$

for all (a,l) in the support of the invariant distribution given s, then a higher capital income tax $\tau^{**} > \tau^*$ reduces the steady-state capital-labour ratio. If (14) holds in reverse, then a higher capital income tax increases the steady-state capital-labour ratio.

Like the previous result, Proposition 6 relies on Lemma 3 to sign the direct response following an increase in the capital income tax rate on savings, and then exploits our main results, Theorems 1–3. We emphasize once more that this proposition is a fairly direct application of our methods, but the economics is both interesting and non-trivial. In particular, there is (and can be) no prediction about the saving behaviour of *all* households, some of which can and often will go against the aggregate. Additionally, (14) is a local condition and applies only at current prices given the initial capital tax τ^*. Hence, our approach, eschewing strong monotonicity requirements and focusing on aggregate behaviour and local conditions, is critical for this result.[28]

The proposition shows how under a simple condition (summarized in equation (14)), we can make sure that standard neoclassical comparative statics hold in the presence of self-control and temptation considerations. At the same time, as in Proposition 3 for the systematic misperceptions case, this result also highlights that when the relevant condition is reversed, standard neoclassical comparative statics can be easily overturned. Such reverse comparative statics do not require extreme parameters and can hold under reasonable economic conditions.[29]

We next explain the logic of condition (14), further clarifying when the reverse comparative static result holds. First, condition (14) does not have any direct or distributive effects, because tax proceeds are not transferred back to households and do not affect the marginal utility of consumption (though it is easy to generalize this condition to the case in which there are such rebates). Second, the intuition should be understood in terms of the effects of capital income taxes on the marginal utility of consumption and marginal temptation costs. To explain this in the clearest possible way, let us ignore uncertainty (assuming that l takes a single value) and again assume full depreciation. Let us also define the shorthand $\hat{R} = (1+R)/(1-\tau) = f'(k^*)$

28. The statement and proof of the proposition also exploit the fact that all households have positive assets. This follows from Proposition 5.

29. For example, suppose v has a positive third derivative, which implies that $v''(c(s(a,l),l')) - v''(\theta Rs(a,l) + wl' - \underline{a}) < 0$ in the numerator on the left-hand side of (14). If u', v', and u'' are all uniformly bounded from below, and v''' is bounded from below by a large enough positive constant, then the reverse comparative static will hold (recall that l has bounded support so that the propensity to save out of assets, $s(a,l)/a$, is uniformly bounded from above).

to denote the (pre-tax) market price of capital. In this case, (14) becomes:[30]

$$(1-\tau^*)\hat{R}\frac{v''(y^*)}{v'(y^*)} \geq \left((1-\tau^*)\hat{R}-1\right)\frac{u''(c^*)+v''(c^*)}{u'(c^*)+v'(c^*)}, \tag{15}$$

where c^* is steady-state consumption and y^* is steady-state wealth. Since $(1-\tau^*)\hat{R} > (1-\tau^*)\theta\hat{R} > 1$ and $u+v$ is strictly concave, this condition immediately implies that when temptation utility is convex, (15) always holds and thus neoclassical comparative statics generalize readily to models with self-control and temptation. Conversely, however, reverse comparative statics apply when v is strictly concave. For example, when $v=u$, constant or increasing absolute rate of risk aversion is sufficient to reverse (15) and thus the standard comparative statics.

The economic intuition for comparative static reversals is also interesting. As opposed to the standard utility function u, the temptation utility v can be concave or convex, even though Gul and Pesendorfer (2004) focus on the case where v is convex. Concavity in this case would imply that households have an incentive to smooth their wealth, since a smoother wealth profile lowers temptation costs. If this wealth smoothing motive is sufficiently strong—in particular, stronger than the consumption smoothing motive—a lower capital income tax rate encourages lower savings in order to achieve a smoother wealth profile.

Finally, we will use the self-control and temptation preferences to show how our methods can be applied for deriving new distributional comparative statics (see Jensen, 2018 for more on distributional comparative statics). To do this in the simplest possible way, we combine these preferences with misperceptions about future labour endowments, which helps us isolate the effects of self-control considerations (abstracting from other effects of changes in labour endowments).[31]

Specifically, all households believe future endowments are given by distribution μ^*, and without loss of any generality, we suppose that initially $\mu^*=\mu$, where μ denotes the true distribution of labour endowments. We then consider a mean-preserving spread of μ^* to μ^{**}. The generalized Euler equation for this case is

$$\begin{aligned}
&-u'((1+R)a+wl-s(a,l))-v'((1+R)a+wl-s(a,l))\\
&\quad+\max\Bigg\{\delta(1+R)\int u'((1+R)s(a,l)+wl'-s(s(a,l),l'))\\
&\quad+v'((1+R)s(a,l)+wl'-s(s(a,l),l'))\\
&\quad-v'((1+R)s(a,l)+wl'-\underline{a})\mu(dl'),u'((1+R)a+wl-\underline{a})\\
&\quad+v'((1+R)a+wl-\underline{a})\Bigg\}=0
\end{aligned}$$

30. In deterministic models, direct responses are always determined because one can apply the implicit function theorem (IFT) to (steady state) Euler equations. In the steady state of a representative household economy, $k^*=a^*=s(a^*;w,R)$. When $T_t=\tau\hat{R}k_t$, steady-state consumption and wealth are thus $c^*=((1-\tau)\hat{R}-1)a^*+wl+\tau\hat{R}k^*=(\hat{R}-1)a^*+wl$ and $y^*=(1-\tau)\hat{R}a^*+wl+\tau\hat{R}k^*=\hat{R}a^*+wl$. The (steady state) Euler equation is therefore $\left(1-\frac{1}{\delta(1-\tau)\hat{R}}\right)\cdot\{u'((\hat{R}-1)a^*+wl)+v'((\hat{R}-1)a^*+wl)\}-v'(\hat{R}a^*+wl)=0$. Applying the IFT to determine $da^*/d\tau$, and using that the Euler equation must hold, one obtains $da^*/d\tau\geq 0 \Leftrightarrow$ (15).

31. We also note that such misperceptions may be quite natural in general, because estimating future distributions is difficult for many households.

Although determining the effects of such distributional shifts is in general challenging, by Lemma 3 comparative statics simply turn on whether the integrand in the GEE is convex or concave as a function of l'. Denoting the initial wealth distribution and consumption function by μ_y^* and c, respectively, convexity therefore holds when

$$u'''(c(y)) + v'''(c(y)) - v'''(y - \underline{a}) \geq 0 \text{ for } y \in \text{Supp}(\mu_y^*) \,. \tag{16}$$

Proposition 7. *A mean-preserving spread of perceived future labour endowments increases the steady-state capital-labour ratio when (16) holds, and reduces the steady-state capital-labour ratio when this inequality is reversed.*

We are once again unaware of any similar results in the literature. Nevertheless, this proposition is intuitive. Consider first the benchmark case with no temptation costs ($v = 0$). In this case, (16) reduces to the well-known "prudence" condition for precautionary savings, and the mean-preserving spread increases savings when consumers are prudent and find it optimal to raise their precautionary savings.

In the presence of temptation utility, there are additional effects. First, if v has a negative fourth derivative, we obtain a prudence effect working through temptation costs, reinforcing the precautionary savings effect (when v has a negative fourth derivative, $v'''(c(y)) - v'''(y - \underline{a}) \geq 0$ in (16)). In contrast, when v has a positive fourth derivative, then temptation considerations work against precautionary savings. The economic intuition is again related to wealth smoothing: a smoother consumption profile implies a more varied wealth profile. When there are strong wealth smoothing motives, soaring consumption becomes costly and discourages precautionary savings, potentially reversing standard comparative statics.

Like in Propositions 3 and 6, the current result shows how local conditions on the quantitative balance between competing effects determine whether a given change in environment leads to higher or lower steady-state capital-labour ratio. In this case, intuitively, wealth smoothing motives can easily reverse prudence effects when the latter are bounded. For example, when $u'''(c) \leq A$ for some $A > 0$ and for all c, and $v''''(c) \geq B > 0$, the following condition is sufficient to reverse (16): $B > -A/\underline{a}$ where $\underline{a} < 0$ is the borrowing limit.[32] This condition crystallizes the intuition that distributional comparative statics with self-control and temptation preferences depend on whether consumption smoothing or wealth smoothing is more important.

4.4. *Quasi-hyperbolic preferences*

We now study the applications of our methods to quasi-hyperbolic preferences, studied among others by Phelps and Pollak (1968), Laibson (1997), Barro (1999), Harris and Laibson (2001), Krusell *et al.* (2002), Balbus *et al.* (2015), and Laibson *et al.* (2020). Despite the popularity and the broad range of applications of these preferences, comparative static analysis is even more challenging in this case, because of dynamic inconsistency.[33] The literature distinguishes

32. Since $y \geq c(y)$ and v has a positive fourth derivative, $u'''(c(y)) + v'''(c(y)) - v'''(y - \underline{a}) \leq u'''(c(y)) + v'''(c(y)) - v'''(c(y) - \underline{a}) = u'''(c(y)) + \int_{c(y)-\underline{a}}^{c(y)} v''''(\tau)d\tau \leq A + \underline{a}B < 0$. This channel can be referred to as "temperance," capturing the aversion to fluctuations in wealth. (16) illustrates the tension between consumption smoothing working through $v(c(y))$ versus wealth smoothing encapsulated in $-v(y - \underline{a})$. We can also note that a positive fourth derivative is necessary and sufficient for the latter effect to dampen precautionary savings (put differently, $v'''(c(y)) < v'''(y - \underline{a})$ if and only if v has a positive fourth derivative).

33. The exception is for the deterministic logarithmic utility case, which is observationally equivalent to the dynamically consistent and fully rational benchmark, as noted in Barro (1999) and Krusell *et al.* (2002).

between the naive and sophisticated versions of hyperbolic discounting. In the former, households do not recognize that they will change their plans in the future, while in the latter they do and thus understand that they are playing a game with their future selves (and savings are determined by the time-stationary Markovian equilibria of this dynamic game). Both cases can be studied with the approaches proposed in this paper and satisfy the assumptions in Lemma 1. In the sophisticated case, the household saving decisions in (8) solve:

$$s(a,l) \in \arg\max_{y\in B((1+R)a+wl)} u((1+R)a+wl-y)+\beta\delta\int W(y,l')\,\mu(dl'),$$

where $B((1+R)a+wl)=\{y\in[\underline{a},\overline{a}]: y\leq(1+R)a+wl\}$, and the continuation utility is given as

$$W(y,l)=u((1+R)y+wl-s(y,l;w,R))+\delta\int W(s(y,l),l')\,\mu(dl'). \tag{17}$$

In the naive case, on the other hand, current selves believe, incorrectly, that future selves will discount geometrically with the "long-run" discount factor δ.[34] Here, the correctly anticipated future saving function $s(y,l)$ in (17) is replaced with a misperceived saving function $s^f(y,l)$ determined as in the benchmark neoclassical consumption-savings problem:

$$s^f(a,l)\in\arg\max_q u((1+R)y+wl-q)+\delta\int V(q,l')\,\mu(dl'),$$

where $V(y,l)=\max_q u((1+R)y+wl-q)+\delta\int V(q,l')\;\mu(dl')$. The naivety in this formulation is rooted in the fact that anticipated future savings will persistently differ from actual future savings. The next proposition establishes comparative static results for the naive quasi-hyperbolic model that parallels the results in Section 4.2.

Proposition 8. *Assume that a fraction $\alpha\in[0,1]$ of households are naive quasi-hyperbolic with $\beta<1$ and long-run discount factor $\delta<1$, and the remaining fraction $1-\alpha$ of households are rational with discount factor $\delta<1$. Then an increase in α or a decrease in either β or δ reduces the steady-state capital-labour ratio.*

An immediate implication of the proposition is that the steady-state capital-labour ratio will be lower in an economy with naive quasi-hyperbolic households than in the benchmark neoclassical model (assuming the same long-run discount factor). As all of our previous results, this result naturally generalizes to the more realistic situation where (measurable) subsets of households may have different discount factors and/or utility functions. This proposition once again exploits Lemma 3 and then applies our main theorems, and the proof in online Appendix C (supplementary material) clarifies that the mathematical arguments are analogous to the ones in Section 4.2 and the effects of capital income taxes with naive quasi-hyperbolic preferences can be studied in the same way. We omit these results to avoid repetition.

Comparative statics in the sophisticated case are more challenging, however, because of the strategic interactions between different selves. This can be seen from the generalized Euler equation in terms of time-stationary savings, $s(a,l)=y$ (where we are again assuming full

34. This is similar to the misperception about future discount factors in Section 4.2. See also the proof of Proposition 8.

depreciation and ignoring the borrowing limit):

$$
\begin{aligned}
&- u'((1+R)a + wl - y) \\
&\quad + \underbrace{\delta(1+R)\int u'((1+R)y + wl' - s(y,l'))\mu(dl')}_{\text{“Standard Impact”}} \\
&\quad \underbrace{- \delta(1-\beta)\int \frac{\partial u((1+R)y + wl' - s(y,l')))}{\partial y}\mu(dl')}_{\text{“Future Selves Adjustment”}} = 0.
\end{aligned}
$$

If we define L as in Lemma 3 and impose Assumption 2, we can establish similar results to those for the naive model. However, inspection of the previous equation reveals that part *3* of Assumption 2 (Monotonicity in future savings) might not hold due to the impact of the “Future Selves Adjustment,” especially when β is “low” and consequently the conflict between current and future selves is severe.

In this case, the comparative statics with sophisticated quasi-hyperbolic households is more complex and potentially more interesting. In the next subsection, we show how numerical analysis can be blended with our methods to make progress in this case. Here, as a final result, we provide a basic intuition for the types of results that arise in this case by focusing on the special case where uncertainty about endowments is very small, which leads to approximately no precautionary savings. First, let us follow Harris and Laibson (2001) and write the Generalized Euler Equation in terms of consumption C as:

$$
u'(C(y_t)) = (1-\tau)\hat{R}\delta \int \left(1 + (\beta - 1)C'(y_{t+1})\right) u'(C(y_{t+1}))\mu(dl_{t+1}), \tag{18}
$$

where $\tau \in [0,1)$ is the capital income tax, $\hat{R} = (1+R)/(1-\tau)$ is again the (pre-tax) market price of capital, $y_t = (1-\tau)\hat{R}a_t + wl_t + T_t$ is current wealth, $y_{t+1} = (1-\tau)\hat{R}[y_t - C(y_t)] + wl_{t+1} + T_{t+1}$ is next period’s wealth, and $C'(y)$ denotes the derivative of the consumption function with respect to current wealth, y. Next assuming CRRA utility with rate of risk-aversion γ, we show in online Appendix C (supplementary material) that as uncertainty about future labour endowments vanishes, the Generalized Euler equation converges to

$$
((1 - C'(y))\hat{R})^{\gamma} = (1-\tau)\hat{R}\delta\left(1 + (\beta - 1)C'(y)\right), \tag{19}
$$

where y is the steady-state level of wealth in the limit economy with no labour endowment uncertainty. Even though we cannot apply Lemma 3, we can use the implicit function theorem to conclude that a lower capital income tax will increase savings *if and only if* (see online Appendix C (supplementary material) for details):

$$
(1 - C'(y))^{-1} > \frac{1-\beta}{\beta}\frac{1-\gamma}{\gamma}. \tag{20}
$$

Conversely, when this condition is reversed, comparative statics of capital income taxes are also reversed.

A couple of additional observations are useful. First, when $\beta \to 1$, there is limited conflict between current and future selves and standard comparative statics apply. Similarly, these comparative static results also hold when $\gamma \geq 1$ (including the logarithmic utility case). In contrast, when $\beta \to 0$ and $\gamma < 1$, the right-hand side diverges and this condition is violated (the left-hand side is bounded above by $1+\beta^{-1}(((1-\tau)\delta)^{-1}-1)$, and thus $\gamma < (1-\tau)(1-\beta)\delta$ is sufficient to ensure the reverse inequality). Condition (20) and this discussion also provide an intuition about why lower capital taxes can reduce long-run capital-labour ratio: as (18) shows, a higher marginal utility in the future will be associated with a higher marginal utility today, which implies that more savings in the future will go together with more savings today. This linkage will be particularly strong when γ is low (high intertemporal elasticity) and when β is low (which in turn implies that $(1-\tau)\delta(1-\beta)$ is large). These strategic interactions between selves makes current consumption very sensitive to future savings, and consequently, greater future savings induced by lower taxes will lead to even more savings today. Because this cannot be sustained in steady state, the steady-state level of savings will have to decrease.

4.5. *Numerical analysis for robust comparative statics*

In this subsection, we show how our general methods can be combined with numerical analysis to obtain additional results and insights. We establish, in particular, that it is possible to blend our Theorems 1–3 with numerical analysis of how aggregate savings respond to changes in the environment at given prices. This will enable us to establish how steady-state equilibria of economies with complex behavioural biases respond to changes in environment, without having to compute a new general equilibrium and the associated asset distributions. Moreover, we will see that this theoretical-cum-numerical analysis can sometimes be conducted without even having to compute the initial steady-state distributions.

For concreteness, we focus on the sophisticated version of the quasi-hyperbolic model introduced in the previous subsection (where a fraction α are sophisticated quasi-hyperbolic and the rest are fully rational). Let us first consider the effects of an increase in α starting from the fully rational benchmark with $\alpha = 0$ to complete the implications of sophisticated quasi-hyperbolic discounting on steady-state equilibria and illustrate how numerical analysis can be used in the context of our approach.

The steady-state equilibrium and asset distribution are straightforward to compute numerically in the benchmark Bewley–Aiyagari model with time-separable preferences, geometric discounting, random labour endowments and borrowing limits. We continue to focus on *i.i.d.* labour endowments and specialize the economy to the case where the labour endowment can take either a "low" or a "high" value (see the notes to Figure 4 for further details).

In our baseline numerical exercise, we follow Krueger *et al.* (2016) and choose $k^* = 5.3$ (which is approximately the average non-housing wealth of US households in 2006, excluding the top 1%) and a capital share of 0.36. We then calibrate A in the usual way by targeting the capital-output ratio (see Krueger *et al.*, 2016, p. 865). This leads to steady-state factor prices of $R^* = R(k^*) = 0.04$ and $w^* = w(k^*) = 1.36$. We can then compute numerically the partial-equilibrium response of households to a reduction in β, which is equivalent to comparing the neoclassical benchmark economy to an economy where a subset $\alpha \in (0,1)$ of households have "sophisticated" quasi-hyperbolic preferences. Note that we cannot apply Lemma 3 to determine the partial-equilibrium response in this case, because the Euler equation does not provide a global characterization of savings behaviour. In fact, as shown by the vertical segments in the figure, the time-stationary saving function is discontinuous.

Once this numerical analysis determines whether aggregate savings at these prices increase or decrease, we can apply Theorems 1–3 and derive the direction of change for the new steady state. Figure 4 summarizes the main idea: all we need to do is to numerically compute the shift of the aggregate savings schedule, and the rest of the work is done by our main theorems. Specifically, Panel A of this figure shows the saving functions of rational and sophisticated

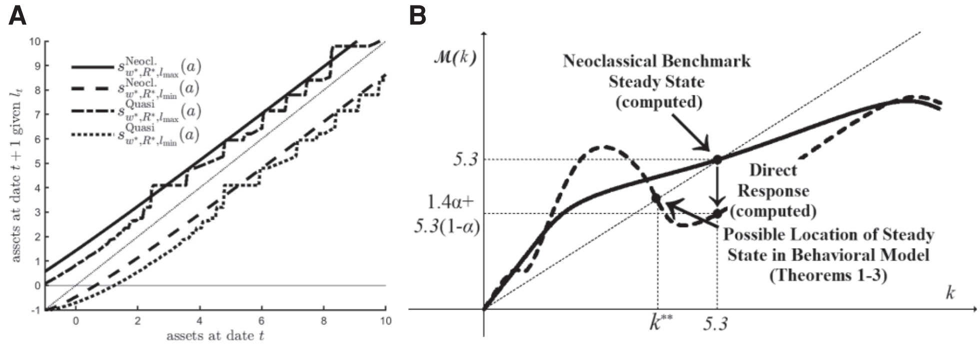

Saving Functions of Neoclassical and Sophisticated Quasi-Hyperbolic Households

Introducing Quasi-hyperbolic Households Reduces Aggregate Savings

FIGURE 4

Change in Aggregate Savings and Equilibrium Adjustment to the Introduction of Sophisticated Quasi-Hyperbolic Households *Panel A*: Benchmark simulation: $k^* = 5.3$, $R^* = R(k^*) = 0.04$, $w^* = w(k^*) = 1.36$. $\underline{a} = -1$, $\delta = 0.947$, $u(c) = c^{0.5}$, $l_{min} = 0.1$, $l_{max} = 1.9$, $p_{min} = p_{max} = 0.5$, $f(k) = 1.163k^{0.36}$, $\Delta = 0.1$. Behavioural households: As in benchmark simulation except $\beta = 0.96$. Vertical segments in the saving functions represent discontinuities. *Panel B*: aggregate savings in the benchmark steady-state simulation is 5.3. Average savings of behavioural households given R^* and w^* is 1.4. With a fraction $\alpha \in (0, 1)$ of behavioural households, the direct effect is thus $1.4\alpha + 5.3(1 - \alpha) - 5.3 < 0$. Computational Notes: Benchmark model computed with the IID Aiyagari EGP algorithm of Kaplan (2017). Quasi-hyperbolic case computed with Ego Loss algorithm of Jensen (2022)

quasi-hyperbolic households for the best and the worst realizations of labour endowments. We can see the jumps in consumption in the quasi-hyperbolic case, which accord with the results in Harris and Laibson (2001). Once we know the saving functions, we can compute the steady-state asset distributions and the aggregate savings levels depicted in Panel B. The solid curve depicts the neoclassical benchmark, while the dashed curve shows the same economy when a fraction $\alpha \in (0, 1)$ of households have sophisticated quasi-hyperbolic preferences. The figure demonstrates that there is a negative direct response at $k^* = 5.3$ from the introduction of quasi-hyperbolic households. The figure also confirms that there are no downward jumps as guaranteed by Lemma 1 and, consequently, the post-tax steady-state capital-labour ratio must be at a point like k^{**}.

When,as in Figure 4, the counterfactual experiment introduces behavioural biases in an otherwise neoclassical economy, our numerical analysis can be further simplified by skipping the computation of the initial steady-state distributions entirely, because it only uses information on steady-state factor prices. In the most common approach to quantitative analysis, the researcher targets some aggregate quantities (such as the aggregate capital-labour ratio, the capital-output ratio or the interest rate), which then pin down steady-state prices. For our computational step, all we need are these steady-state prices, and once these are determined, we can readily move to the partial-equilibrium step of determining whether aggregate savings following the change in environment increase at these prices.

Our next application is more involved. In this case, we start from the steady state of a sophisticated quasi-hyperbolic economy with $\beta < 1$. We then consider an increase in the capital income tax rate τ and numerically study its (direct) impact on aggregate savings. Once this direct response is obtained, we again apply Theorems 1–3.

In this exercise, we distinguish two cases, both depicted in Figure 5. In the first (Panels A and B), we choose a "high" short-run discount factor, $\beta = 0.94$. In the second (Panels C and D),

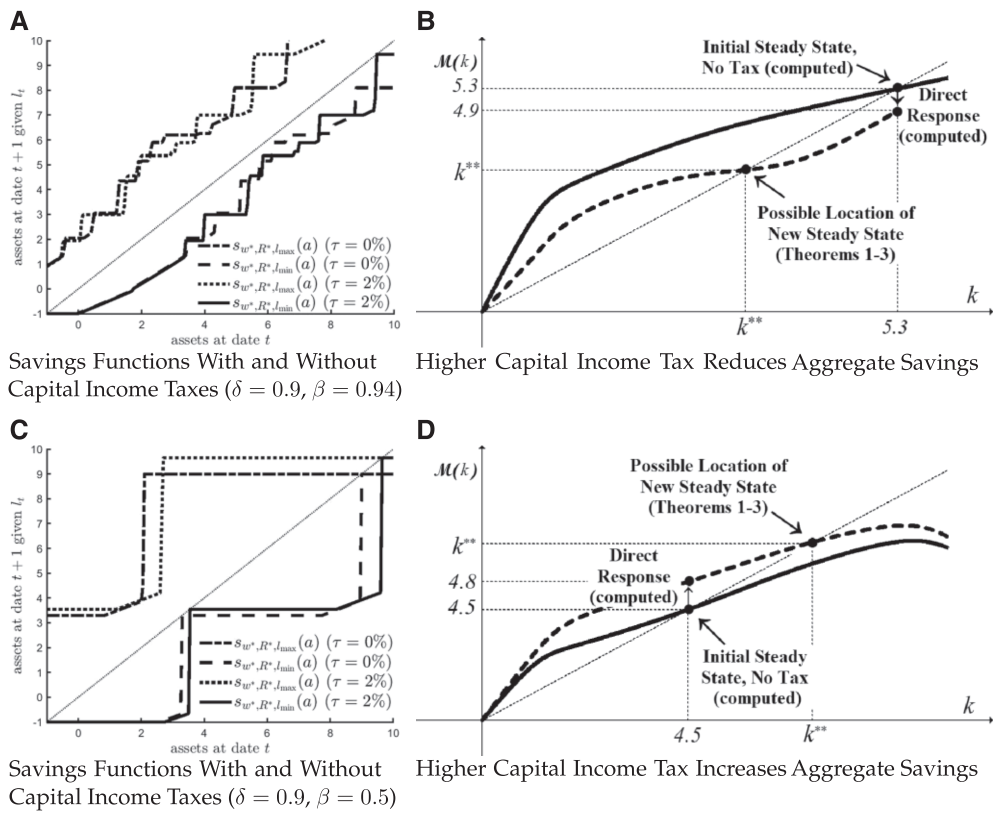

FIGURE 5

Comparative Statics in the "Sophisticated" Quasi-Hyperbolic Model *Panel A*: $k^* = 5.3$, $R^* = 0.19$ (pre-tax), $w^* = 2.30$, $\tau = 0, 0.02$. $\underline{a} = -1$, $\delta = 0.9$, $\beta = 0.94$, $u(c) = c^{0.1}$, $l_{min} = 1$, $l_{max} = 3$, $p_{min} = p_{max} = 0.5$, $f(k) = 1.97k^{0.4}$, $\Delta = 0.1$, $T = \tau k$ (the capital share is from Kaplan). *Panel B*: Aggregate savings in the steady-state simulation with $\tau = 0$ is 5.3. Aggregate savings when $\tau = 2\%$ is 4.9. Direct Response is $4.9 - 5.3 < 0$. *Panel C*: $k^* = 4.5$, $R^* = 0.84$ (pre-tax), $w^* = 6.36$, $\tau = 0, 0.02$. $\underline{a} = -1$, $\delta = 0.9$, $\beta = 0.5$, $u(c) = c^{0.5}$, $l_{min} = 1$, $l_{max} = 3$, $p_{min} = p_{max} = 0.5$, $f(k) = 5.81k^{0.4}$, $\Delta = 0.1$, $T = \tau k$. *Panel D*: Aggregate savings in the steady-state simulation with $\tau = 0$ is 4.5. Aggregate savings when $\tau = 2\%$ is 4.8. Direct Response is $4.8 - 4.5 > 0$. *No transfers in panel D*: If $T = 0$, aggregate savings when $\tau = 2\%$ is 4.77, hence direct response is positive whether or not taxes are transferred to households. *Computational notes*: Computed with the Ego Loss algorithm of Jensen (2022)

we choose a "low" discount factor, $\beta = 0.5$, which generates a strong conflict between current and future selves. Since this low discount factor makes households much less willing to save, we can no longer achieve the same aggregate capital-labour ratio target with the parameter choices described previously. In this application, therefore, we depart from our baseline in two ways. First, we relax the capital-labour ratio target to 4.5. Though this number is different from our benchmark and the baseline of Krueger *et al.* (2016), it is in the ballpark of the recent US-wealth-to-GDP numbers (*e.g.* FRED, 2023). Second, we follow Kaplan (2017) and set the capital share in the production function to 0.4 (instead of the 0.36 used in Figure 4).[35]

35. The case in Panels C and D has some parallel to the illustrative result we presented in the previous section when Condition (20) is violated (for an economy where uncertainty about endowments becomes very small). Crucially, however, in the current case endowment uncertainty is not "small" and as in the previous application, the Euler equation does not provide a global characterization of savings behaviour (so, again, Lemma 3 cannot be used).

Panels B and D of Figure 5 show that in these two cases, aggregate savings move in opposite directions (in Panel B the initial solid curve is above the after-tax dashed curve, while in Panel D it is the other way around). In particular, in the first case, higher capital income taxes reduce aggregate savings at given factor prices (as in the neoclassical benchmark). Then Theorems 1–3 ensure that the steady-state capital-labour ratio decreases (Panel B). In contrast, the second case illustrates how this standard comparative statics result can be reversed in the sophisticated quasi-hyperbolic model. Specifically, in Panel D, the direction of the direct response is reversed and the higher capital income tax raises aggregate savings at given prices. Then from Theorems 1–3, the new steady state must have higher capital-labour ratio. This discussion also reveals how the same combination of numerical analysis and our theorems can be applied in other settings.

5. CONCLUDING REMARKS AND FUTURE DIRECTIONS

A common conjecture is that equilibrium analysis becomes excessively challenging in the presence of behavioural preferences and biases, thus implicitly justifying a focus on models with time-additive, dynamically consistent preferences and rational expectations. In this paper, we demonstrated that, in the context of one-sector behavioural growth models, this conjecture is not necessarily correct. Results concerning the direction of change in the long run (or "robust comparative statics" for the steady-state equilibrium) can be obtained for a wide range of behavioural preferences and rich heterogeneity. Put simply, our main results state the following: for any change in policy or underlying production or preference parameters of the model, we first determine whether at the initial capital-labour ratio (or at the initial pre-tax/distortion vector of prices) aggregate savings increases or decreases; this step involves no equilibrium analysis, but only the determination of what the average of individual optimization decisions given prices is. Critically, this needs to be done only at a single vector of prices (or at a single capital-labour ratio), because our condition is completely "local." Then under fairly mild regularity conditions (satisfied for all behavioural preferences we have discussed in this paper), no matter how complex the equilibrium responses are, they will not overturn the direction of the initial change and thus the steady-state equilibrium will involve a greater capital-labour ratio (and the changes in prices that this brings). Conversely, if the initial change is a decline in aggregate savings at the initial capital-labour ratio, the long-run capital-labour ratio will decline.

At the root of this result is a simple and intuitive observation: in the one-sector model, the only way the direction of the initial impetus can be reversed is by having the equilibrium response to this initial shock to go strongly in the opposite direction. For example, savings could decline strongly in response to a higher capital-labour ratio. But either such an equilibrium response would still not overturn the initial increase in aggregate savings, in which case the conclusion about the steady-state equilibrium applies. Or it would overturn it and reduce the long-run capital-labour ratio, but in this case the perverse effect would go in the direction of strengthening, not reversing, the initial increase in savings (since it was the *higher* capital-labour ratio that induced the decline and aggregate savings).

We illustrated these comparative statics by working through one-sector growth models embedding three different types of behavioural considerations: (1) systematic misperceptions; (2) self-control and temptation preferences; and (3) naive and sophisticated quasi-hyperbolic discounting. In all three cases, we showed that our approach can be applied relatively straightforwardly and leads to results that are, to the best of our knowledge, new in the literature. We also identified conditions under which these behavioural biases reverse standard neoclassical comparative static results (*e.g.* with respect to declines in capital income taxes). In each case, this reversal takes place along the lines of our main result: behavioural preferences change the direction of the direct response and this initial impetus then leads to a change in the same direction in the long-run equilibrium.

We further showed how our key results can be blended with numerical analysis. In particular, in the context of sophisticated quasi-hyperbolic model, we showed how simple numerical analysis can be used to sign the direction of (partial equilibrium) responses, which can then be used to determine the full general equilibrium comparative statics. We believe this combination of new theory and numerical analysis can be used in other settings as well.

Our analysis has several limitations, which point to interesting areas for future research. As already implied by our discussion, there are several important cases in which our results do not apply. First, with more general preferences than those considered in this paper, the upper hemi-continuity of the saving correspondence established in Lemma 1 may no longer apply and we cannot rule out the case in Panel D of Figure 1. Second, non-stationary belief formation processes, whereby partial-equilibrium and long-run responses are driven by very different beliefs, would also render our theorems inapplicable. Finally and most importantly, our results do not apply when there are multiple aggregate state variables rather than the single state variable as in our (one-sector) behavioural growth model. In such cases, as is well known from other comparative static settings, even shifts that lead to positive responses for each dimension can induce negative overall effects because of cross-dimension dependencies (or because of failure of negative semi-definiteness of local Jacobians even when they have negative diagonal elements). With multiple aggregates, similar results would necessitate at least some supermodularity conditions for the set of state variables. Acemoglu and Jensen (2015) provide some conditions for comparative statics in neoclassical economies with two aggregates, and developing such results in the richer setting we consider here is one future direction for research.

Another evident limitation of our approach bears repeating at this point: our focus has been on comparative statics, and thus on qualitative rather than quantitative results. Many questions in modern macroeconomics necessitate quantitative analysis, and the quantitative impact of a policy change may critically depend on behavioural biases and the exact structure of preferences even if the direction of long-run change does not. An obvious but challenging area for future research is to investigate when certain quantitative conclusions may not depend on certain types of behavioural biases or heterogeneity (*e.g.* in the sense that as behavioural assumptions are modified, quantitative change in some key variables remains near changes implied by a benchmark model).

Perhaps the most important area for future research is to extend the analysis to non-steady-state environments. Behavioural considerations may matter greatly for the response of an economy to recessionary shocks, and overoptimism and other misperceptions may be important during temporary periods of rapid expansion. In principle, one could study whether different behavioural biases change the direction of response to various macroeconomic changes, such as interest rate cuts, but this is challenging because these biases will also alter the future evolution of state variables. One approach may be to leverage the fact that, in some cases, the impact on future variables will be small relative to current effects, though there may also be other fruitful approaches, and we leave the exploration of these issues to future work.

APPENDIX

A. Proofs

Proof of Lemma 1. Throughout, the household index i is omitted to simplify notation. Fixing e and the future saving function on the right-hand side of (8), we can write the current self's decision problem as

$$s(y;e) \in \arg \max_{a' \in [\underline{a}, \max\{\overline{a}, y\}]} u^{\epsilon}(y - a') + M(a'),$$

where M is a function that does not depend on current assets a. Note that we here express savings as a function of wealth $y = (1+R)a + wl + T$ (that this is possible can be seen from (8)). The constraint correspondence and the first term in the objective function are clearly continuous in y, the latter because $u^{\epsilon}(\cdot)$ is continuous by assumption. To see that M is continuous in a' independently of e and the future savings function used to define M, write it out in full: $M(a') = \int V^{\epsilon^1}(y - \tilde{s}(y;e_1), (1+R_2)\tilde{s}(y;e_1) + w_2 l_2 + T_2 - \tilde{s}((1+R_2)\tilde{s}(y;e_1) + w_2 l_2 + T_2; e_2), \ldots) P(de'|e)$, where $\tilde{s} \in \mathcal{S}$ and $\mathcal{S}$ is the space of uniformly bounded measurable functions with the weak* topology.[36] Then rewrite as follows:

$$\begin{aligned}
(M(a')) &= \int\int V^{\epsilon^1}(y - \tilde{s}(y;e_1), (1+R_2)\tilde{s}(y;e_1) + w_2 l_2 + T_2 - \tilde{s}((1+R_2)\tilde{s}(y;e_1) \\
&\quad + w_2 l_2 + T_2; e_2), \ldots) Q(dl_1|e, e'\backslash\{l_1\}) P(d(e'\backslash\{l_1\}|e) \\
&= \int\int V^{\epsilon^1}(y - \tilde{s}(y;e_1), (1+R_2)\tilde{s}(y;e_1) + w_2 l_2 + T_2 \\
&\quad - \tilde{s}((1+R_2)\tilde{s}(y;e_1) + w_2 l_2 + T_2; e_2), \ldots) f(l_1|e, e'\backslash\{l_1\}) dl_1 P(d(e'\backslash\{l_1\}|e) \\
&= \int\int V^{\epsilon^1}(y - \tilde{s}(y;e_1), (1+R_2)\tilde{s}(y;e_1) + w_2 l_2 + T_2 \\
&\quad - \tilde{s}((1+R_2)\tilde{s}(y;e_1) + w_2 l_2 + T_2; e_2), \ldots) \eta(d(y, l_1)|e, e'\backslash\{l_1\})) P(d(e'\backslash\{l_1\}|e)
\end{aligned}$$

where Q and P is a disintegration family of measures for the projection $e' \mapsto l_1$ (*e.g.* see Chang and Pollard, 1997), f is the (continuous) Radon–Nikodym derivative of $Q(\cdot|e, e'\backslash\{l_1\})$ with respect to the Lebesgue measure on $\mathbb{R}_+$, and $\eta(A \times B|e, e'\backslash\{l_1\}) = \int\int 1_{l_1}(A) 1_{(1+R_1)a'+w_1 l_1+T_1}(B) f(dl_1|e, e'\backslash\{l_1\}) dl_1$ (which is continuous in a'). Note that the disintegration exists because P is a Radon measure.[37] Because there is no loss of generality in assuming that V is bounded by an integrable function (see the discussion immediately prior to Assumption 1), it follows by the dominated convergence theorem that $M(a';e)$ is continuous in a'. We remark that this argument is essentially the "change of variable" argument used in Harris and Laibson (2001)'s proof of existence and continuous dependence on the short-run discount factor in the quasi-hyperbolic model (see *e.g.* Lemma 5 in that paper), except we consider a more general measure space. Because V is continuous, it is also clear that if $\tilde{s}_n \to s$ in $\mathcal{S}$, then $M(a';e,\tilde{s}_n) \to M(a';e,\tilde{s})$. Finally, we also have continuity in e (because e, unlike a' and $\tilde{s}$, does not directly enter the integrand, $P(\cdot|e)$ is continuous in e by assumption, and the integrand is bounded). By the theorem of the maximum, the arg max, $F(y;e,\tilde{s}) = \arg\max u^{\epsilon}(y - a') + M(a')$, is non-empty and upper hemi-continuous in y, e, and $\tilde{s}$. Because the objective function is continuous in y, and has strictly increasing differences in (y, a') if and only if $u^{\epsilon}(\cdot)$ is strictly concave, it follows from Topkis' theorem that any solution $s(y;e)$ must be increasing in y (see Topkis, 1978). This is true, *inter alia*, if $\tilde{s} = s$, that is, if the future savings function we fix to begin with is equal to s. Since $y = (1+R)a + wl + T$, we conclude that any TSSF must be increasing in assets. Therefore, if the time-stationary savings correspondence

36. The original topology is the essential supremum norm topology. The weak-$*$ topology $\sigma(X', X)$ on the set of savings functions is then defined for the dual pair (X, X'), where X' is the topological dual of X.

37. The set of bounded sequences with the supremum norm is a separable metric space, hence any Borel probability measure is a Radon measure.

$S(a;e)$ has well-defined compact values, the least and greatest selections must be increasing in assets.

Next, note that when we fix $\tilde{s}$, the set of selections from $G(\tilde{s}) = F(\cdot;\cdot,\tilde{s})$ is a singleton in any (quotient) space of measurable functions $s:(y,e)\mapsto s(y,e)$ under almost everywhere equal equivalence. This is because for any fixed e, as we have just proved, any selection must be increasing—hence any two selections can differ only at points where both are discontinuous (which implies they are equal at all but an at most countable number of points given any e).[38] Because G is upper hemi-continuous on $\mathcal{S}$, it is also continuous on $\mathcal{S}$. Since $\mathcal{S}$ is convex, and compact by the Banach-Alaoglu theorem, existence of a fixed point (a TSSF) now follows from the Schauder–Tychonoff fixed point theorem. The time-stationary savings correspondence is the set of fixed points, hence compactness of $S(a,e)$ and its upper hemi-continuous dependence on a and e now follow from a standard argument (*e.g.* see the last paragraph of the next proof). □

Proof of Lemma 2. To shorten expressions, we set $z^i=(l^i,\epsilon^i)\in Z^i=[l^i_{\min},l^i_{\max}]\times E^i$. We first show that $k\in\mathcal{M}^\theta(k)$ if and only if k is a steady-state equilibrium. Consider $m(k)\in\mathcal{M}^\theta(k)$. Using that $c^i(a^i;z^i,w,R,T^i)=(1+R)a^i+wl^i+T^i-s^i(a^i;z^i,w,R,T^i)$ and that $R(k)=(1-\tau(k))f'(k)-\Delta$, $w(k)=(1-\omega(k))(f(k)-f'(k)k)$, $\tau(k)f'(k)k+\omega(k)(f(k)-f'(k)k)=\int T^i di+G(k)$, and $\int\hat{l}^i\ \ di=1$ we have $m(k)=f(k)+(1-\Delta)k-G(k)-(1+(1-\tau(k))f'(k_t)-\Delta)k-(1-\omega(k))(f(k)-f'(k)k)-\int T^i\ \ di\ +\int s^i$ $(\frac{\hat{a}^i(k)}{\int\hat{a}^i(k)\,di}k,\hat{z}_i,w(k),R(k),T^i)\ di=\int s^i(\frac{\hat{a}^i(k)}{\int\hat{a}^i(k)\,di}k,\hat{z}_i,w(k),R(k),T^i)\ di$, where for all i, s^i is some selection from S^i. Because $s^i(\hat{a}^i(k)/(\int\hat{a}^i(k)\ di)k,\hat{z}_i,w(k),R(k),T^i)$ has the image measure of $(\hat{a}^i(k),\hat{z}^i(k))$ under $(a^i,z^i)\mapsto s^i(a^i/(\int\hat{a}^i(k)\ di)k,z^i,w(k),R(k),T^i)$, it has distribution $\lambda^i(\cdot,Z^i)$ where $\lambda^i=\lambda^i(k;s)$ was defined just prior to Definition 3.[39] Hence $m(k)=\int\hat{a}^i(k;s)\ di$. It follows that if $m(k)=k$, then $k=\int\hat{a}^i(k;s)\ di$, that is, the capital market must clear. Further, $\lambda^i(k;s)$ must then be an invariant distribution for (*a.e.*) i. Because prices satisfy (1) and (2) by construction, we thus have $k=m(k)\in\mathcal{M}^\theta(k)$ whenever k is a steady-state equilibrium. Conversely, if k is a steady-state equilibrium, then there exist invariant distributions λ^i and $\lambda^i(k;s)=\lambda^i$ where $k=\int\hat{a}^i(k;s)\ di$, hence $k=m(k)\in\mathcal{M}^\theta(k)$ when $m(k)$ corresponds to the family of selections s.

Next, we show that $\mathcal{M}^\theta$ has convex values. Here, we use the definition that a correspondence $F:A\to 2^B$ is type I (type II) monotone if $a\succeq\tilde{a}$ and $b\in F(a)$ ($\tilde{b}\in F(\tilde{a})$) implies the existence of $\tilde{b}\in F(\tilde{a})$ ($b\in F(a)$) such that $\tilde{b}\succeq b$. For a selection $s^i\in S^i$ write $\lambda^i(k,K;s)$ if $\lambda^i(A\times B)=\int_{A\times l^i,\epsilon^i}q^i(B,l^i,\epsilon^i)1_A(s^i(a^i\frac{k}{K};l^i,\epsilon^i,w,R,T^i))\lambda^i(da^i,dl^i,\epsilon^i)$ (for all Borel sets). Note that the right-hand side defines an adjoint Markov operator $r^i_{s^i,k,K}$ in the usual way, and the savings correspondence S^i thus defines an adjoint Markov correspondence which we denote by $T^i_{k,K}=\{r^i_{s^i,k,K}:s^i\in S^i\text{ and is measurable}\}$. Also, let $\mathcal{A}^i_k(K)\equiv\{\hat{a}^i\sim\lambda^i(\cdot,Z^i)\in\mathcal{P}([\underline{a}^i,\overline{a}^i]):\lambda^i\in T^i_{k,K}\lambda^i\}$ denote stationary assets. If $\lambda^i\in T^i_{k,K}\lambda^i$, then $\lambda^i\in\Omega(\mu_z)=\Omega(\mu_z)=\{\lambda^i\in\mathcal{P}([\underline{a}^i,\overline{a}^i]\times Z^i):\lambda^i([\underline{a}^i,\overline{a}^i],B)=\mu_z(B)\}$ since $\lambda([\underline{a}^i,\overline{a}^i],B)=\int q(z^i,B)[\int_{[\underline{a}^i,\overline{a}^i]}\lambda(da^i|z^i)]\mu_z(dz^i)=\int q(z^i,B)\mu_z(dz^i)=\mu_z(B)$. By our Lemma 1 and Theorem B1 in Acemoglu and Jensen (2015), $T^i_{k,K}:\Omega(\mu_z)\to\Omega(\mu_z)$ is (weak-$*$) upper hemi-continuous in a^i and K, and Type I and Type II monotone in the

38. Let A be the set of discontinuities. Then the (Lebesgue) measure is $\int\int 1_y(A_e)dy\mu(de)$ where μ is the (product) Lebesgue measure on $\mathbb{R}^5$ and $A_e=\{y:s(\cdot,e)\text{ is discontinuous at }y\}$. Clearly $\int 1_y(A_e)da=0$ if A_e is at most countable.

39. $\text{Prob}(s^i(\hat{a}^i(k)/(\int\hat{a}^i(k)di)k,\hat{z}_i,w(k),R(k),T^i)\in A)=\int 1_{s^i(a^i(k)/(\int\hat{a}^i(k)\,di)k,z_i,w(k),R(k),T^i)}(A)\lambda^i(da^i,dz^i)=\lambda^i(A,Z^i)$.

order $\succeq_{A-FOD}$ defined by $\lambda \succeq_{A-FOD} \tilde{\lambda} \Leftrightarrow [\lambda(\cdot, B) \succeq_{FOD} \tilde{\lambda}(\cdot, B)$ for all $B \in \mathcal{B}(Z)]$. It is also increasing in that order in any parameter for which $s^i((a^i/K)k, z^i, w(k), R(k), T^i)$ is increasing. By Theorem 3 in that paper and our Lemma 1, the set of fixed points is type I and II monotone in K^{-1}. Moreover, by Theorem B3 in Acemoglu and Jensen (2015), it is non-empty and upper hemi-continuous in K. By the definition of $\succeq_{A-FOD}$ it follows that $\mathcal{A}_k^i(K)$ is non-empty and decreasing in K. By Richter's theorem (see Aumann, 1965), $\int \mathcal{A}_k^i(\cdot)\ di$ is convex-valued, and by Theoresupposed tom 4 in Acemoglu and Jensen (2015), it has decreasing least and greatest selections. A convex and real-valued correspondence whose least and greatest selections are decreasing must have a convex set of fixed points (this statement is straight-forwardly verified graphically). We conclude that $\mathcal{M}^\theta(k) = \{K : K \in \int \mathcal{A}_k^{\theta,i}(K)\ di\}$ is convex.

To see that the market correspondence $\mathcal{M}^\theta(k) = \{K : K \in \int \mathcal{A}_k^{\theta,i}(K)\ \ di\}$ is upper hemi-continuous, note that its graph is $\{(k, K) : (K, k, K) \in \text{Graph}[\int \mathcal{A}_k^{\theta,i}(K)\ \ di]\}$ where $\text{Graph}[\int \mathcal{A}_k^{\theta,i}(K)\ di] = \{(K, k, Z) : Z \in \int \mathcal{A}_k^{\theta,i}(K)\ di\}$ is a closed set since $\int \mathcal{A}_k^{\theta,i}(K)\ di$ is upper hemi-continuous in k and K. That $\mathcal{M}^\theta(k)$ is compact follows now from boundedness (savings correspondences have compact ranges). Finally, $\mathcal{M}^\theta(k)$ begins above the 45 ° line and ends below it. The former is obvious since $f(0) = 0$ and therefore $\mathcal{M}^\theta(0) = \{0\}$. The latter is true since consumption is non-negative, hence $\mathcal{M}^\theta(k) \leq f(k)$, and for sufficiently large k, $f(k) \leq k$ because the production technology is effectively compact. □

Proof of Theorem 1. The proof relies on the following lemma (whose proof is similar to Lemma 2 and relegated to online Appendix C (supplementary material)).

Lemma 4 (Mean asset holdings and shifts in the market correspondence)**.** *Assume that households satisfy the conditions of Lemma 1, and let $k^* \in \mathcal{M}^{\theta^*}(k^*)$ be either the least steady state* $\inf\{k : k \in \mathcal{M}^{\theta^*}(k)\}$ *or the greatest steady state* $\sup\{k : k \in \mathcal{M}^{\theta^*}(k)\}$ *given an environment $\theta^* \in \Theta$. Consider a different environment $\theta^{**} \in \Theta$. Then the population's mean asset holdings increase (decrease) at k^* when the environment changes from θ^* to θ^{**} if and only if the market correspondence "shifts up" ("shifts down") at k^*, i.e. provided there exists $\tilde{k} \in \mathcal{M}^{\theta^{**}}(k^*)$ with $\tilde{k} \geq k^*$ ($\tilde{k} \leq k^*$).*

We provide the proof for the case in which the market correspondence shifts up (the down case is analogous).

Sufficiency: By Lemma 4, there exists $\tilde{k} \in \mathcal{M}^{\theta^{**}}(k^*)$ with $\tilde{k} \geq k^*$. Since $\mathcal{M}^{\theta^{**}}$ ends below the 45° (Lemma 2), it must begin above and end below the 45° line on the interval $[k^*, +\infty)$. $\mathcal{M}^{\theta^{**}}$ is also upper hemi-continuous and convex valued (Lemma 2, again), hence it intersects the 45° line at some $k^{**} \in [k^*, +\infty)$. This yields a steady-state equilibrium $k^{**} \geq k^*$ given environment θ^{**}, and by assumption, this is the unique steady-state equilibrium.[40]

Necessity: Assume that $k^{**} \geq k^*$ and that the change from θ^* to θ^{**} does not increase the households' mean asset holdings. By Lemma 4, the market correspondence then does not shift up at k^*. So $\sup \mathcal{M}^{\theta^{**}}(k^*) < k^*$ since the market correspondence is closed. But then since the market correspondence ends below the 45° line and is upper hemi-continuous and convex valued, $\mathcal{M}^{\theta^{**}}$ must then intersect with the 45° at least twice on the interval $[k^*, +\infty)$. This contradicts that the economy has a unique non-trivial steady state given θ^{**}. □

Proof of Theorem 2. Since the market correspondence is compact-valued, a sufficiently small change in the environment can lead to existing equilibria disappearing but not to the creation

40. The same conclusion follows by instead considering a single-valued market correspondence that is continuous but for jumps up (see Appendix B).

of new equilibria. In particular, no new equilibrium can be created below the least equilibrium which must therefore increase by the argument used to prove Theorem 1. This argument obviously also applies to the greatest equilibrium; and in both cases necessity follows by the argument from Theorem 1 as well. □

Proof of Theorem 3. Let k^* denote the greatest steady state. Repeating the argument used to prove the "sufficiency" part of Theorem 1, $\mathcal{M}^{\theta^{**}}$ must have a fixed point on $[k^*, +\infty)$. The result for the least steady state is proved analogously. □

Proof of Proposition 1. We suppress the transfers T^i to simplify notation. Throughout, prices are fixed at the levels determined by the initial capital-labour ratio k^*. Both rational and behavioural households use current discount factor δ and correctly anticipate the saving function of future selves, $s(a; l, \epsilon)$, where ϵ encapsulates beliefs about future discount factors. The only difference between the two types of households is that rational ones believe (correctly) that $\hat{P}(\epsilon_\tau = \delta) = 1$ for all $\tau > t$, where $\hat{P}$ is ϵ_τ's marginal belief distribution, while behavioural ones believe (incorrectly) that $\hat{P}(\epsilon_\tau = \hat{\delta}) = 1$ for all $\tau > t$. We focus here on the case where $\hat{\delta} > \delta$. For both rational and behavioural households, (8) implies the Euler equation $-u'((1+R)a + wl - y_t) + \max\{\epsilon(1+R)\int u'((1+R)y_t + wl' - s(y_t; l', \epsilon'))\mu(dl')\hat{P}(d\epsilon'), u'((1+R)a + wl - \bar{a})\} = 0$, where $(a, l, \epsilon) \in [\underline{a}, \overline{a}] \times [l_{\min}, l_{\max}] \times \{\delta, \hat{\delta}\}$, and $y_t = s(a; l, \epsilon)$ is the solution ($s(y_t; l', \epsilon')$ is time-stationary savings of a "future self" with labour endowment l' and discount factor ϵ'). Since rational households have $\epsilon = \delta$ and $\hat{P}(\epsilon' = \delta) = 1$, the Euler equation reduces in this case to the benchmark Euler equation $-u'((1+R)a + wl - y_t) + \max\{\delta R \int u'((1+R)y_t + wl' - s^{\text{Neocl.}}(y_t; l, \delta))\mu(dl'), u'((1+R)a + wl - \bar{a})\} = 0$ where $s^{\text{Neocl.}}(a; l, \delta) = y_t$ denotes the rational households' saving function conditioned on $\epsilon = \delta$ (since rational households place zero probability on $\epsilon \neq \delta$, we do not need to specify savings when $\epsilon \neq \delta$ in this case).

Next, let $s^{\text{Beh.}}(\cdot; l, \hat{\delta})$ denote the solution to the Euler equation when $\epsilon = \hat{\delta}$ and $\hat{P}(\epsilon' = \hat{\delta}) = 1$. Clearly $s^{\text{Beh.}}(\cdot; l, \hat{\delta}) = s^{\text{Neocl.}}(\cdot; l, \hat{\delta})$; that is, $s^{\text{Beh.}}(\cdot; l, \hat{\delta})$ solves the benchmark Euler equation with $\hat{\delta}$ in place of δ. Since this equation has a unique solution and its the left-hand side is increasing in δ, it follows immediately from Lemma 3 that $s^{\text{Beh.}}(a; l, \hat{\delta}) \geq s^{\text{Neocl.}}(a; l, \delta)$ for all l and a, that is, behavioural households anticipate greater savings in the future than rational households (for all a and l). The behavioural households' time-stationary saving function given $\epsilon = \delta$ is the solution to the Euler equation with $\epsilon = \delta$ and $\hat{P}(\epsilon' = \hat{\delta}) = 1$. Equivalently, $s^{\text{Beh.}}(a; l, \delta) = y_t$ where y_t must solve $-u'((1+R)a + wl - y_t) + \max\{\delta(1+R)\int u'((1+R)y_t + wl' - s^{\text{Beh.}}(y_t; l', \hat{\delta}))\mu(dl'), u'((1+R)a + wl - \bar{a})\} = 0$. Comparing with the benchmark Euler equation above and using (i) that $s^{\text{Beh.}}(y_t; l', \hat{\delta}) \geq s^{\text{Neocl.}}(y_t; l, \delta)$ for all l' and y_t, and (ii) that u' is decreasing, it follows from a second application of Lemma 3 that $s^{\text{Beh.}}(a; l, \delta) \geq s^{\text{Neocl.}}(a; l, \delta)$ for all a and l. Since consumption at date t is $c_t = Ra + wl - s^{\text{Beh.}}(a; l, \delta)$, the budget constraint necessarily holds (the dynamic inconsistency is embedded in the beliefs).

Since behavioural households (just like rational households) will always "observe" $\epsilon = \delta$ at the current date, $s^{\text{Beh.}}(a; l, \delta)$ is the TSSF which the behavioural households will actually adopt at every date. As mentioned before, $s^{\text{Neocl.}}(a; l, \delta)$ is the TSSF which rational households will adopt at every date. Thus at every date, behavioural households will save more than rational households (on average). Taking as environment the fraction of behavioural households in the population, $\theta := \alpha \in [0, 1] = \Theta$, raising α therefore entails a positive direct response for the subset of rational households that is interchanged with behavioural households. Since the saving function of the remaining (rational) households is not impacted by such an increase in α, the (average) direct response of Definition 5 is positive. The conclusion of the proposition now follows from one of Theorems 1–3. The proof for the case where $\hat{\delta} < \delta$ is analogous and may be omitted (in this case, $s^{\text{Beh.}}(a; l, \delta) \leq s^{\text{Neocl.}}(a; l, \delta)$ for all a and l above and we proceed as a moment ago except that the direct response will now be negative). □

B. Changes in the environment: a topological approach, discussion of related literature

Since this section's observations may be of independent interest and apply not only to market correspondences, we are going to view the market correspondence $\mathcal{M} : K \times \Theta \to 2^{\mathbb{R}}$, $K \subseteq \mathbb{R}$, more abstractly and impose any necessary assumptions directly. Denote by $m_S^\theta(k) = \inf \mathcal{M}^\theta(k)$ and $m_L^\theta(k) = \sup \mathcal{M}^\theta(k)$ the least and greatest selections, and by $k_S^\theta = \inf\{k \in K : k \in \mathcal{M}^\theta(k)\}$ and $k_L^\theta = \sup\{k \in K : k \in \mathcal{M}^\theta(k)\}$ the least and greatest fixed points (when they exist, which of course they do if $\mathcal{M}$ is a market correspondence). Now equip Θ with an order as well as a topology (in the simplest situation where we consider a change in just a single parameter, Θ may be taken to be a subset of $\mathbb{R}$, and these would therefore be the usual/Euclidean order and topology, respectively). We also introduce some additional terminology: A function $m : \Theta \to \mathbb{R}$ is (i) *increasing* if $\theta \leq \hat{\theta} \Rightarrow m(\theta) \leq m(\hat{\theta})$ for all $\theta, \hat{\theta} \in \Theta$, and (ii) *locally increasing at* $\theta^* \in \Theta$ if $\theta \leq \hat{\theta} \Rightarrow m(\theta) \leq m(\hat{\theta})$ for all $\theta, \hat{\theta}$ in an open neighbourhood of θ^*. Finally, $\mathcal{M}$ *begins above and ends below the 45° line* if $m_*(\inf K, \theta) \geq \inf K$ and $m^*(\sup K, \theta) \leq \sup K$. The following is proved in online Appendix C (supplementary material) where we also present a corollary that directly addresses one-sector growth models from the topological perspective .

Theorem 4 (Abstract shifts in fixed point correspondences)**.** *Consider an upper hemi-continuous and convex valued correspondence* $\mathcal{M} : K \times \Theta \to 2^{\mathbb{R}}$ *where K is a compact subset of* $\mathbb{R}$ *and* Θ *is a compact subset of an ordered topological space. Suppose that the graph begins above and ends below the 45° line for all* $\theta \in \Theta$*. Then the least and greatest fixed points* k_S^θ *and* k_L^θ *are increasing in* θ *if for all* $\theta^* \in \Theta$, $m_L^\theta(k_L^{\theta^*})$ *and* $m_S^\theta(k_S^{\theta^*})$ *are locally increasing in* θ *at* θ^*.

Note that in all cases, "curve shifting theorems" such as Theorem 4 can be used in our setting because (i) Lemma 2 has established the requisite properties of the market correspondence; and (ii) Lemma 4 allows us to relate increases in mean savings/assets with "shifts up" in the market correspondence.

Most of the results in the literature are similar to Corollary 2 in Milgrom and Roberts (1994) which shows that when the equivalent of our market correspondence $\mathcal{M}$ is "continuous but for jumps up" and its graph shifts up (meaning that $m_L^\theta(k)$ and $m_S^\theta(k)$ are increasing in θ *for all* k), then the least and the greatest fixed points increase.[41] Let us refer to this well-known result as the "for all k curve shifting theorem." The key thing to note is that since the "curve" must shift up for all k (for all capital-labour ratios in our setting), it requires information not only about how savings change for the prices determined in the original steady state; it requires that we have such information for (all) capital-labour ratios/prices. Both Acemoglu and Jensen (2015) and Light and Weintraub (2021) define "local positive shocks" as changes in parameters that increase savings for all capital-labour ratios.[42] In conventional settings with rational expectations, such requirements can be imposed, even if they are quite demanding. When the economic problems involve rich and variegated behavioural preferences and biases, they become essentially untenable. It is against this background that Theorem 4 should be evaluated. It shows that

41. $\mathcal{M}$ is *continuous but for jumps up* if it has convex values, $\limsup_{x^n \uparrow x} m^*(x^n, t) \leq m_L^\theta(k)$, and $\liminf_{x^n \downarrow x} m_*(x^n, t) \geq m_S^\theta(k)$. Acemoglu and Jensen (2013) proves that if $\mathcal{M}$ is upper hemi-continuous in k and has convex values, then it is continuous but for jumps up.

42. Note that (ii) above integrates seamlessly with the approach in Light and Weintraub (2021), hence both our main results and curve shifting arguments are easily integrated with the mean-field games literature.

if $\mathcal{M}$ is upper hemi-continuous in (k,θ) (rather than just in k, *cfr.* footnote 41), the same conclusion requires only that the correspondence shifts up at the least and the greatest fixed points, k_S^θ and k_L^θ. The results presented in Section 3 similarly require only local shifts in steady states. That we only need to verify that $\mathcal{M}$ shifts up *locally*, in particular, *at* the steady states, enables us to separate direct responses (or the "all-else-equal" behaviour) from equilibrium responses.

To explain a little further, let us consider a particularly simple case where a dynamic economy can be reduced to a fundamental equation of the form

$$G(k_t,k_{t-1},\theta)=0, \tag{21}$$

where $\theta\in\mathbb{R}$ is an exogenous parameter, $k_t\in\mathbb{R}$ is capital, or the capital-labour ratio, at date t and $G:\mathbb{R}^3\to\mathbb{R}$ a smooth function. In this case, the market correspondence can be defined as

$$\mathcal{M}^\theta(k)=\{\hat{k}:G(\hat{k},k,\theta)=0\}\,. \tag{22}$$

In the Ramsey–Cass–Koopmans model, for example, $G(k_t,k_{t-1},\theta)=0\Leftrightarrow k_t=g(k_{t-1},\theta)$, and then $\mathcal{M}^\theta(k)=g(k,\theta)$. Clearly, k^* is a steady state given θ if and only if $k^*\in\mathcal{M}^\theta(k^*)$. Note, however, that (21)—even in the more general form $0\in G(k_t,k_{t-1},\theta)$ where G is a correspondence—is not general enough to nest our one-sector behavioural growth model (because we also need to condition on the distribution of assets). Nevertheless, (21) is useful to provide the technical intuition for our main results since both in the case of (22) and our Definition 3, the market correspondence is constructed by conditioning on the information that the capital-labour ratio in question, k, has to be consistent with a steady-state equilibrium. In particular, the fact that, with the conditioning on the steady state k^*, (22) a one-dimensional fixed point problem allows us to use "curve shifting" arguments without imposing any type of monotonicity on the dynamical system defined by (21) (see also Acemoglu and Jensen, 2015 for a related discussion of non-monotone methods). Given $\mathcal{M}^\theta(k)$ and this construction, Theorem 4 and the results presented in Section 3 enable us to predict how the greatest and the least steady states vary with θ when $\mathcal{M}^\theta(k)$ shifts up locally starting at these steady states (and provided that $\mathcal{M}$ satisfies the relevant theorem's regularity conditions).

The added generality and flexibility is considerable. In many applications, including the problem of equilibrium analysis in the behavioural growth model we focus on in this paper, the conditions for the "for all k curve shifting theorem" will not hold even if (21) applies. This is for both substantive and technical reasons. Substantively, as already mentioned, in economies such as the one-sector behavioural growth model the possible heterogeneity in the responses of agents to changes in the environment would often preclude such uniform shifts. To see the technical problem, suppose that we were checking these conditions using the implicit function theorem. That would amount to verifying that $\frac{dk}{d\theta}>0$ for all $\tilde{k}$ while $G(k,\tilde{k},\theta)=0$ holds. But since the implicit function theorem requires as a minimum that $D_kG(k,\tilde{k},\theta)\neq 0$, and "running through all $\tilde{k}$'s" will almost invariably violate this condition for some $\tilde{k}$, this method will generally fail (order theoretic methods are of no help here either; and of course, it is *not* enough to show that $\frac{dk}{d\theta}>0$ for almost every $\tilde{k}$ because any point we fail to check may precisely be a point where the market correspondence "jumps"). When we only need to check local conditions, these difficulties are bypassed.

Acknowledgments. We are grateful to the editor Kurt Mitman and three anonymous referees for numerous very useful comments and suggestions. We also thank Xavier Gabaix for several very useful conversations. Additional thanks are due to Drew Fudenberg, Marcus Hagedorn, David Laibson, Paul Milgrom, Kevin Reffett, and Christian Wolf, as well as participants at the TUS-IV conference in Paris, and seminar participants at Lund University and the University of Oslo for helpful comments and suggestions. Acemoglu gratefully acknowledges financial support from the Army Research Office.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

Data Availability Statement

The paper does not use original data. The code underlying this research is available on Zenodo at https://dx.doi.org/10.5281/zenodo.7695696.

REFERENCES

ACEMOGLU, D. (2009), *Introduction to Modern Economic Growth* (Princeton, New Jersey: Princeton University Press).

ACEMOGLU, D. and JENSEN, M. K. (2010), "Robust Comparative Statics in Large Static Games", *IEEE Proceedings on Decision and Control*, **49**, 3133–3139.

ACEMOGLU, D. and JENSEN, M. K. (2013), "Aggregate Comparative Statics", *Games and Economic Behavior*, **81**, 27–49.

ACEMOGLU, D. and JENSEN, M. K. (2015), "Robust Comparative Statics in Large Dynamic Economies", *Journal of Political Economy*, **123**, 587–640.

ACHDOU, Y., HAN, J., LASRY, J.-M., *et al.* (2022), "Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach", *Review of Economic Studies*, **89**, 45–86.

AHN, H., KAPLAN, G. and MOLL, B., *et al.* (2018), "When Inequality Matters for Macro and Macro Matters for Inequality", *NBER Macroeconomics Annual*, **32**, 1–75.

AIYAGARI, S. R. (1994), "Uninsured Idiosyncratic Risk and Aggregate Saving", *Quarterly Journal of Economics*, **109**, 659–684.

AIYAGARI, S. R. and MCGRATTAN, E. R. (1998), "The Optimum Quantity of Debt", *Journal of Monetary Economics*, **42**, 447–469.

AL-NAJJAR, N. I. (2004), "Aggregation and the Law of Large Numbers in Large Economies", *Games and Economic Behavior*, **47**, 1–35.

AUBIN, J.-P. and FRANKOWSKA, H. (1990), *Set-Valued Analysis* (Boston: Birkhauser).

AUMANN, R. J. (1965), "Integrals of Set-Valued Functions", *Journal of Mathematical Analysis and Applications*, **12**, 1–12.

BACKUS, D. K., ROUTLEDGE, B. R. and ZIN, S. E. (2004), "Exotic Preferences for Macroeconomists", *NBER Macroeconomics Annual*, **19**, 319–390.

BALBUS, L., REFFETT, K. and WOZNY, L. (2015), "Time Consistent Markov Policies in Dynamic Economies with Quasi-Hyperbolic Consumers", *International Journal of Game Theory*, **44**, 83–112.

BARRO, R. J. (1999), "Ramsey Meets Laibson in the Neoclassical Growth Model", *Quarterly Journal of Economics*, **114**, 1125–1152.

BECKER, G. S. (1962), "Irrational Behavior and Economic Theory", *Journal of Political Economy*, **70**, 1–13.

BENJAMIN, D. J, RABIN, M. and RAYMOND, C. (2015), "A Model of Nonbelief in the Law of Large Numbers", *Journal of the European Economic Association*, **14**, 515–544.

Board of Governors of the Federal Reserve System (US) (2023), "Households and Nonprofit Organizations; Net Worth, Level [TNWBSHNO]", retrieved from FRED (Federal Reserve Bank of St. Louis, https://fred.stlouisfed.org/series/TNWBSHNO).

CAO, D. and WERNING, I. (2018), "Saving and Dissaving with Hyperbolic Discounting", *Econometrica*, **86**, 805–857.

CHANG, J. T. and POLLARD, D. (1997), "Conditioning as Disintegration", *Statistica Neerlandica*, **51**, 287–317.

DEATON, A. (1991), "Savings and Liquidity Constraints ", *Econometrica*, **59**, 1221–1248.

EPSTEIN, L. G. and HYNES, J. A. (1983), "The Rate of Time Preference and Dynamic Economic Analysis", *Journal of Political Economy*, **91**, 611–635.

EPSTEIN, L. G. and ZIN, S. E. (1989), "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework", *Econometrica*, **57**, 937–969.

EPSTEIN, L. G. and ZIN, S. E. (1991), "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis", *Journal of Political Economy*, **99**, 263–286.

FUDENBERG, D. and LEVINE, D. K. (2006), "A Duel Self Model of Impulse Control", *American Economic Review*, **96**, 1449–1476.

FUDENBERG, D. and LEVINE, D. K. (2012), "Timing and Self-Control", *Econometrica*, **80**, 1–42.

GABAIX, X. (2014), "A Sparcity Based Model of Bounded Rationality", *Quarterly Journal of Economics*, **129**, 1661–1710.

GABAIX, X. (2017), "Behavioral Macroeconomics via Sparse Dynamic Programming" (Working Paper, Harvard).

GILBOA, I. (1987), "Expected Utility with Purely Subjective Non-Additive Probabilities", *Journal of Mathematical Economics*, **16**, 65–88.

GILBOA, I. and SCHMEIDLER, D. (1995), "Case-Based Decision Theory", *Quarterly Journal of Economics*, **110**, 605–639.

GUL, F. and PESENDORFER, W. (2001), "Temptation and Self-Control", *Econometrica*, **69**, 1403–143.

GUL, F. and PESENDORFER, W. (2004), "Self-Control and the Theory of Consumption", *Econometrica*, **72**, 119–158.

HARRIS, C. and LAIBSON, D. (2001), "Dynamic Choices of Hyperbolic Consumers", *Econometrica*, **69**, 935–957.
JENSEN, M. K. (2018), "Distributional Comparative Statics", *Review of Economic Studies*, **85**, 581–610.
JENSEN, M. K. (2022), "The Ego Loss Approach to Dynamic Inconsistency" (Working Paper, University of Nottingham).
KAPLAN, G. (2017), "IID Aiyagari EGP Algorithm (Matlab)". Available at http://benjaminmoll.com/ha˙codes/.
KOOPMANS, T. C. (1960), "Stationary Ordinal Utility and Impatience", *Econometrica*, **28**, 287–309.
KREPS, D. M. and PORTEUS, L. (1978), "Temporal Resolution of Uncertainty and Dynamic Choice Theory", *Econometrica*, **46**, 185–200.
KRUEGER, D., MITMAN, K. and PERRI, F. (2016), "Macroeconomics and Household Heterogeneity", in Uhlig, H. and Taylor, J. (eds) *Handbook of Macroeconomics*, Vol 2B, Chapter 11 (Amsterdam: Elsevier) 843–921.
KRUSELL, P., KURUSCU, B. and SMITH, A. A. (2002), "Equilibrium Welfare and Government Policy with Quasi-Geometric Discounting", *Journal of Economic Theory*, **105**, 42–72.
KRUSELL, P., KURUSCU, B. and SMITH, A. A. (2010), "Temptation and Taxation", *Econometrica*, **78**, 2063–2084.
KRUSELL, P. and SMITH, A. A. (2003), "Consumption–Savings Decisions with Quasi-Geometric Discounting", *Econometrica*, **71**, 365–375.
LAIBSON, D. (1997), "Golden Eggs and Hyperbolic Discounting", *Quarterly Journal of Economics*, **112**, 443–477.
LAIBSON, D., MAXTED, P. and MOLL, B. (2020), "Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy" (Working Paper).
LI, H. and STACHURSKI, J. (2014), "Solving the Income Fluctuation Problem with Unbounded Rewards", *Journal of Economic Dynamics and Control*, **45**, 353–365.
LIGHT, B. (2020), "Uniqueness of Equilibrium in a Bewley-Aiyagari Model", *Economic Theory*, **69**, 435–450.
LIGHT, B. and WEINTRAUB, G. Y. (2021), "Mean Field Equilibrium: Uniqueness, Existence, and Comparative Statics", *Operations Research*, **70**, 585–605.
MACHINA, M. (1982), "Expected Utility Analysis Without the Independence Axiom", *Econometrica*, **50**, 277–323.
MCFADDEN, D. (1974), "Conditional Logit Analysis of Qualitative Choice Behavior, in Zarembka, P. (ed) *Frontiers in Econometrics* (New York: Academic Press) 105–142.
MILGROM, P. (1994), "Comparing Optima: Do Simplifying Assumptions Affect Conclusions?", *Journal of Political Economy*, **102**, 607–615.
MILGROM, P. and ROBERTS, J. (1994), "Comparing Equilibria", *American Economic Review*, **84**, 441–459.
MILGROM, P. and SHANNON, C. (1994), "Monotone Comparative Statics", *Econometrica*, **62**, 157–180.
PHELPS, E. S. and POLLAK, R. A. (1968), "On Second-Best National Saving and Game Equilibrium Growth", *Review of Economic Studies*, **35**, 185–199.
QUAH, J. K.-H. (2007), "The Comparative Statics of Constrained Optimization Problems", *Econometrica*, **75**, 401–431.
STROTZ, R. H. (1956), "Myopia and Inconsistency in Dynamic Utility Maximization", *Review of Economic Studies*, **23**, 165–180.
SUN, Y. (2006), "The Exact Law of Large Numbers via Fubini Extension and Characterization of Insurable Risks", *Journal of Economic Theory*, **126**, 31–69.
TOPKIS, D. M. (1978), "Minimizing a Submodular Function on a Lattice", *Operations Research*, **26**, 305–321.
UHLIG, H. (1996), "A Law of Large Numbers for Large Economies", *Economic Theory*, **8**, 41–50.
VIVES, X. (1990), "Nash Equilibrium with Strategic Complementarities", *Journal of Mathematical Economics*, **19**, 305–321.