

Liberating Preferences from Balanced Growth: A Case for Birth-Level Normalization

José-Víctor Ríos-Rull^a, Raul Santaella-Llopis^b

^a*University of Pennsylvania, NBER and CEPR*

^b*New York University Abu Dhabi, CEPR and BSE*

Abstract

Balanced growth paths (BGPs) are commonly thought to require constant-relative-risk-aversion (CRRA) preferences and, when leisure is valued, a Frisch elasticity welded to that curvature (King et al., 1988a,b). This restriction, however, is conditional on a paradigm where choices are evaluated in absolute terms, absent an aggregate reference benchmark. Yet, extensive evidence from across psychology, sociology, and other social sciences shows that individual preferences are fundamentally shaped by aggregate benchmarks, in particular those experienced during formative years. Taking this insight to heart, we show that under these preferences, a property we call *pairwise homogeneity* is necessary and sufficient for BGPs, imposing no functional restrictions on the felicity function; standard CRRA is merely a knife-edge sub-manifold. Thus, the traditional CRRA discipline is an artifact of neglecting aggregate benchmarks in preferences, not an inherent property of growth. We extend the analysis to trending hours under *independent pairwise homogeneity*. Implementing this framework via behaviorally grounded birth-level normalization, we show that—anchored on micro-estimates from Blundell et al. (2016)—this single device simultaneously reconciles consumption curvature with the secular decline in hours, breaks the Chetty (2006) asset pricing bound, and rationalizes the cross-cohort transformation of female labor supply. U.S. census microdata on over one million married couples confirm that women’s participation shifts with childhood female-employment environments, validating this framework.

Keywords: Balanced Growth Path, Preferences, Pairwise Homogeneity, Aggregate Benchmarks, CRRA, Frisch

1. Introduction

Macroeconomic models typically require a balanced growth path (BGP). Under the standard assumption of preferences where consumption and leisure are evaluated in absolute terms, without aggregate reference benchmarks, it is well-known that this requirement imposes strict functional form restrictions on the felicity function $u(c_t, h_t)$: preferences must be CRRA in consumption and multiplicatively separable in labor, or log-additive (King et al., 1988a,b). These functional restrictions carry over from infinite-horizon settings to overlapping-generations (OLG) environments and, in recent work by Boppart and Krusell (2020), have been extended to accommodate declining hours, albeit still within the domain of benchmark-independent preferences.

Yet, in contrast to this current paradigm, a rich body of research in psychology, sociology, and other social sciences shows that individual preferences are fundamentally shaped by aggregate reference benchmarks. In particular, this literature argues that individuals evaluate outcomes not in absolute terms, nor against an ever-shifting contemporaneous average, but relative to the standards prevailing during formative years. Intuitively, these early-life benchmarks define the baseline for satisfaction; a cohort born into an era of modern amenities does not perceive them as a source of high absolute utility, but rather as the starting point for evaluating gains—rendering a reversion to previous standards significantly more painful.¹ This implies that as the economy scales, the units of the utility function must scale accordingly. This evidence motivates a broader class of felicity functions: $U(c_t, h_t; C_\tau)$, where c_t and h_t denote individual consumption and labor supply, respectively, and C_τ is the aggregate benchmark representing per-capita consumption at a reference time $\tau \leq t$.

Aggregate benchmarks and BGP. Our main contribution is to characterize which felicity functions within this broader class that explicitly incorporates aggregate benchmarks are consistent with a BGP. We prove that a property we call *pairwise homogeneity* (PH) in the pair (c, C) is both necessary and sufficient for balanced growth. Mathematically, PH is a symmetry property of the felicity function that describes how it responds to the simultaneous scaling of a specific pair of arguments:

$$U(\lambda c_t, h_t; \lambda C_\tau) = \lambda^k U(c_t, h_t; C_\tau)$$

Intuitively, because the pair of arguments are c and C , PH requires that utility depends on the scale of individual consumption relative to its associated benchmark. As the economy grows, the

¹Consider the adoption of modern amenities such as hot showers. While an earlier generation might have perceived the transition from cold to hot water as a substantial gain in absolute utility, a cohort born into an environment where hot water is universal adopts it as a baseline requirement. For this later cohort, the presence of the amenity generates little marginal satisfaction, yet its absence represents a substantial welfare loss. This process reflects a classic paradigm in the social sciences—dating back to Tarde (1903) and formalized in Duesenberry (1949), wherein “luxuries” progressively become behavioral necessities. The notion that reference standards become anchored during formative years is present in Easterlin (1980) and Becker (1996).

aggregate benchmark “lifts” the reference frame of the utility function, ensuring that the trade-offs governing individual choices remain stable even as absolute standards of living rise. The key implication of this result is that PH imposes no restrictions on the functional form of u ; any felicity function satisfying standard regularity conditions is consistent with a BGP, provided it satisfies PH. Consequently, the shape of the utility function is not disciplined by balanced growth.

Sketch of the argument. The logic for PH to deliver consistency with balanced growth is rooted in the stationarity of marginal rates of substitution. Along a BGP, the intertemporal marginal rate of substitution (MRS) must remain constant over time. In the standard paradigm, the constancy of this ratio in the presence of growth forces the marginal utility to be a power function—that is, the CRRA utility. By contrast, under *pairwise homogeneity*, the intertemporal MRS is stationary in all periods independently of the shape of preferences. To see this argument, consider the context of an infinite horizon neoclassical growth model and write the intertemporal MRS between two consecutive periods along the BGP as:

$$\frac{U_c(c_t, h_t; C_\tau)}{\beta U_c(\lambda c_t, h_t; \lambda C_\tau)} = \frac{U_c(c_t, h_t; C_\tau)}{\beta \phi(\lambda) U_c(c_t, h_t; C_\tau)} = \frac{1}{\beta \phi(\lambda)} \quad \forall t$$

where along the BGP the economy grows at constant rate $\lambda - 1$. The first equality states that under *pairwise homogeneity* we can factor out the scaling function $\phi(\lambda)$ from the marginal utility, rendering the intertemporal MRS perfectly stationary along the BGP for *any* regular felicity function. A similar logic applies to the intratemporal margin: *pairwise homogeneity* ensures that the marginal rate of substitution between consumption and labor grows at the rate of the aggregate benchmark, which matches the growth rate of wages. Technically, this occurs because the homogeneity of the utility function implies that the marginal utility of labor inherits a degree of homogeneity exactly one degree higher than that of consumption; consequently, the intratemporal MRS inherits the growth rate of the benchmark. In this way, the rising MRS required to reflect wage growth is absorbed by the aggregate scaling factors rather than the curvature of the utility function. The benchmarks thus perform the structural “lifting” required to offset long-run income and substitution effects, liberating the preferences from specific functional form restrictions. Thus, the standard paradigm (King et al., 1988a,b) occupies a knife-edge sub-manifold of our PH class. In this restricted case, the absence of aggregate benchmarks imposes an additional constraint that forcefully couples the degree of homogeneity to the curvature of the utility function, creating the illusion that BGP consistency is a property of the felicity function’s shape rather than the underlying homogeneity of the environment.

Extension to long-run trends of hours. We extend this characterization to economies with trending hours via *independent pairwise homogeneity* (IPH), where consumption and labor supply (or leisure) are scaled by their respective aggregate benchmarks. This generalized property identifies

the scaling requirements for both margins, nesting the preferences of Boppart and Krusell (2020) as a further restricted sub-manifold. In that setting without aggregate benchmarks, a rigid interlocking of the felicity function’s curvatures is required: the curvature in consumption must be perfectly calibrated to the cross-elasticity and the long-run growth rate of hours. This singular alignment is the necessary cost of neglecting aggregate benchmarks, whereby the functional form itself must mimic the homogeneity of the environment to compensate for the absence of benchmarking. In contrast, in the presence of aggregate benchmarks, our IPH framework delegates the requirement of balanced growth to the property of homogeneity, thereby decoupling the long-run trend from the specific functional form of the utility function.

Implementation via birth-level normalization. We advocate a specific implementation of *pair-wise homogeneity*: birth-level normalization, where τ in an overlappings generation context is set to the birth year. Hence, the aggregate references are the standard of living—and working—prevailing when an agent enters the economy. This choice is grounded in the “impressionable years” hypothesis in psychology (Krosnick and Alwin, 1989), the socialization literature in sociology (Clausen, 1968), and the economics of early-life preference formation (Becker, 1996; Heckman, 2006). Across these disciplines, evidence shows that individuals evaluate lifetime outcomes relative to the societal baselines they were socialized to expect during their upbringing, rather than against absolute levels or ever-shifting contemporaneous averages (Bisin and Verdier, 2001). Empirically, this cohort-based anchor manifests in long-lasting formative shock effects on consumption aspirations (Easterlin, 1980), financial risk-taking (Malmendier and Nagel, 2011), and, most critically, the secular evolution of female labor supply driven by childhood exposure to shifting work norms (Fernández et al., 2004; Fogli and Veldkamp, 2011). Thus, anchoring preferences to this birth-level benchmark is not a mere analytical convenience to decouple growth from the shape of preferences; it is grounded in a realistic representation of how preferences are formed.

Labor margin is central. This decoupling between functional form and balanced growth implies that key preference parameters can be identified from microeconomic, asset-pricing, or business-cycle evidence while remaining fully consistent with long-run macroeconomic trends. In our framework, the consumption benchmark C_τ performs the structural lifting that frees the consumption curvature σ , establishing that balanced growth does not necessitate a CRRA utility shape. In addition, by absorbing the secular growth of real wages, the labor benchmark H_τ ensures that the preference elasticity in hours remains entirely unrestricted. This breaks the parameter interlocking that traditionally welds preferences across margins whenever models move beyond knife-edge log-separability or introduce trending hours. The leisure margin is therefore where the Frisch elasticity, an object of decades of contested estimation, is structurally set free.

A unified view of identification. Anchored on the family-labor-supply estimates of Blundell et al. (2016), we show that our framework does three jobs at once. It reconciles the micro consumption

curvature with the secular decline in hours (Boppart and Krusell, 2020); it frees that curvature for asset pricing, addressing the long-standing tension (Chetty, 2006; Mehra and Prescott, 1985) between the high risk aversion finance demands and the moderate values growth tolerates; and it rationalizes the one first-order fact a benchmark-independent felicity function cannot carry along a balanced growth path, namely, the cross-cohort transformation of married women’s labor supply. The male Frisch margin, stable and well identified, serves as the control case where the two paradigms coincide (Blundell et al., 2016); the household level, and married women’s participation specifically, is where they diverge. Consequently, our empirical strategy focuses primarily on the female margin to measure the role of aggregate benchmarks (e.g., gendered work norms). We bring this last application to U.S. Census microdata on over one million married couples: a wife’s participation rises with the female-employment environment of her own childhood (Fogli and Veldkamp, 2011) and, in the cross-spouse pattern of Fernández et al. (2004), that of her husband’s childhood. While this previous work showed how maternal exposure and localized learning shape women’s long-run employment decisions, their empirical frameworks operate entirely outside the discipline of a balanced growth path. In contrast, our framework embeds these channels through aggregate benchmarks in preferences and consistently with balanced growth. We show evidence of the gendered, early-formed signature for women that birth-level normalization predicts, which survives strict controls that absorb the leading bargaining-based alternatives. Thus, our birth-level normalization is not only theoretically *useful* in reconciling balanced growth with micro-elasticities and asset-pricing facts, but is also explicitly *used* by the data, as our empirical analysis of the cross-cohort transformation of married women’s labor supply shows.

Nesting previous approaches Our framework rationalizes BGP-consistency of several prominent preferences with consumption externalities in the literature (Abel, 1990; Galí, 1994; Campbell and Cochrane, 1999), showing that their capacity to sustain a balanced growth path stems fundamentally from the principle of *pairwise homogeneity*; a consumption externality without this property is not consistent with balanced growth. In the same fashion, we show that ad-hoc scaling “tricks,” such as the time-varying labor disutility weights linked to aggregate technology in Mertens and Ravn (2011), succeed precisely because they implicitly satisfy IPH. Similarly, the recursive preference structure of Jaimovich and Rebelo (2009)—which utilizes an internal, endogenous scaling variable to eliminate wealth effects in the short run while preserving them in the long run—pertains to our IPH class with a specific degree of homogeneity. Our approach turns these previous approaches into preferences with a behavioral foundation by delegating this structural “lifting” to external cohort benchmarks (C_τ and H_τ), which are grounded in psychological and social science evidence on preference formation.

The rest of the paper is organized as follows. Section 2 reviews evidence from psychology, sociology and other social sciences to show why aggregate benchmarks belong in the utility func-

tion, arguing that individual choices are shaped by the standards of our formative years. Section 3 describes the current paradigm: felicity functions without aggregate benchmarks and the functional-form restrictions they impose for BGP consistency. Section 4 introduces our broader class of preferences explicitly incorporating aggregate benchmarks, defines *pairwise homogeneity*, and proves that PH is necessary and sufficient for a BGP with trendless hours, where standard CRRA preferences (as per King et al., 1988a,b) emerge as a knife-edge sub-manifold. Further, we extend our framework to secular trends in hours via *independent pairwise homogeneity*, nesting Boppart and Krusell (2020). Section 5 interprets the degree of homogeneity and links our framework to existing preference structures, showing how *pairwise homogeneity* underlies their capacity to sustain a BGP. Section 6 develops a unified view of identification—the micro anchor, the reconciliation of curvature with the hours decline, the asset-pricing bound, and the gendered census evidence on women’s labor supply. Section 7 collects further implications for cross-sectional and business-cycle evidence. Section 8 concludes.

2. A Case for Birth-Level Normalization

An extensive body of research across psychology, sociology and other social sciences establishes that individual preferences are not evaluated in a vacuum; instead, they are fundamentally shaped by aggregate benchmarks anchored during an agent’s formative years. Rather than reacting strictly to absolute levels, individuals evaluate economic outcomes relative to the baseline standard of living experienced during youth. This perspective captures a classic paradigm in the social sciences, dating back to Tarde (1903) and formalized by Duesenberry (1949), which discusses how expanding aggregate productivity permanently lifts societal expectations, transforming historical luxuries into behavioral necessities. Therefore, grounding preferences in a cohort-specific, birth-level benchmark is not a mere analytical convenience that—under some properties that are the core of our work—decouples growth from preferences; it is a structural representation of how human preferences are formed.

This section reviews the interdisciplinary evidence on benchmarking preferences in formative years across three distinct dimensions. First, Section 2.1 examines the psychological foundations of early-life preference stability. Second, Section 2.2 details the sociological and anthropological channels of socialization and enculturation that transmit these aggregate baselines. Finally, Section 2.3 reviews the economic evidence on cohort-based consumption aspirations and risk-taking, demonstrating how childhood exposure to aggregate environments leaves a permanent signature on lifelong economic choices.

2.1. The “Impressionable Years” and Preference Stability

Across disciplines, the “impressionable years” hypothesis posits that beliefs, values, and attitudes develop at a relatively young age and remain comparatively stable thereafter. In psychology,

the seminal work of Krosnick and Alwin (1989) showcases that late adolescence and early adulthood are periods of peak susceptibility to environmental influence; immediately following this window, susceptibility to behavioral or attitudinal change drops precipitously. This aligns with findings by McCrae and Costa (1994), who emphasize that core personality traits and dispositional tastes stabilize significantly after early adulthood.

The emphasis on the crucial effects of early childhood on adult personality and behavior traces back to the Freudian tradition but finds modern empirical support in developmental psychology (Thompson et al., 2012; Wood et al., 2006). As Becker (1996) notes, “*childhood-acquired habits continue, even though these would not have developed if the environment when growing up had been the same as the environment faced as an adult.*” Becker further argues that this deep-seated influence is entirely consistent with rational choice: “*The Freudian emphasis on the crucial influence of early childhood on later behavior would be consistent with utility-maximizing forward-looking behavior if behavior were highly habitual. For then experiences while a child could have a very large effect on adult preferences and choices*” (Becker, 1996, p. 10). We interpret these insights as a compelling rationale for selecting the environment at birth as the permanent reference unit of account for an individual’s choices.

This stability finds rigorous economic grounding in the work of Heckman (2006) and Cunha and Heckman (2007). Regarding the life cycle of human development, Heckman (2006) argues that both cognitive and non-cognitive traits are highly malleable in early childhood but exhibit significant persistence and diminished plasticity as an individual matures, effectively locking in the reference standards of one’s formative environment. Cunha and Heckman (2007) formalize this through the concept of dynamic complementarity, where early investments in specific orientations and standards increase the productivity of subsequent choices. Thus, birth-level aggregate benchmarks can be interpreted not as an arbitrary static constraint, but a structural reflection of the fact that an individual’s initial calibration to societal norms makes it optimal to maintain those anchors throughout the life cycle.

A leading concrete instance of this formation concerns women’s labor supply. A woman who grows up observing her mother—and, through her, the women of her cohort and place—in paid work forms a different baseline for the consumption–work tradeoff than one who does not, and that baseline persists. This is the mechanism of Fernández et al. (2004) and Fogli and Veldkamp (2011): maternal and local female employment during a cohort’s formative years shape its members’ (and, strikingly, their future spouses’) participation decades later. Crucially, as wartime mobilization made vivid, this is an aggregate rather than an idiosyncratic experience—a cohort-and-place average, exactly the object C_{t-i} (and its labor counterpart H_{t-i}) is meant to capture. Section 6 brings this mechanism to census microdata and shows that it carries the cross-cohort transformation of female participation that a benchmark-independent felicity cannot.

2.2. Socialization and Enculturation

The process of anchoring choices to birth environments finds deep roots in the sociological study of *socialization*. Clausen (1968) provides the seminal framework, defining socialization as the developmental process through which individuals acquire the orientations necessary for functioning within their society. Clausen emphasizes that the “socialization of the infant and child” is the primary mechanism through which societal standards, such as those of consumption and labor, become permanent cognitive anchors. Under this view, our normalization constants, C_{t-i} and H_{t-i} , represent the outcome of the socialization process. This sociological perspective is complemented by the anthropological concept of *enculturation* (Appadurai, 2004). Appadurai (2004) argues that the “capacity to aspire” is itself a cultural fact, constructed within the specific material and social landscape of one’s upbringing. An individual’s horizon of possibilities is not an abstract, universal preference map but is instead scaled to the specific “standard of living” and “norms of effort” prevalent during their entry into society. We formalize the concepts of socialization and enculturation through the birth-level normalization of preferences.

2.3. Cohort-Based Preference Formation

The economic intuition for birth-level normalization finds strong empirical support in the study of consumption aspirations and risk-taking behavior. Easterlin (1980) and Pampel and Peters (1995) show that young adults confronted with a deteriorating economic climate do not simply lower their expectations. As Easterlin (1980) observes: “*Young adults who are confronted with a deteriorating economic climate will not give up the consumption aspirations acquired during their youth.*” To maintain these standards, cohorts often postpone life milestones, such as family formation, rather than adjusting their baseline preferences. This persistent influence of formative years extends beyond consumption to other economic choices such as financial risk-taking. Malmendier and Nagel (2011) show significant birth-year effects in investment behavior, finding that individuals who experience low stock market returns during their youth are less likely to participate in the market or hold equities later in life. These cohort-specific macroeconomic shocks carry a permanent weight in an individual’s perception of risk, mirroring the way consumption benchmarks are anchored at the time of entry into the economy.

In the context of cohort-based preference formation, the literature has focused on two primary theoretical mechanisms: *intergenerational transmission*, where a cohort’s preferences depend on the consumption patterns of the previous generation (Easterlin et al., 1980), and the *intrafamily model*, where preferences are anchored by the standards experienced during childhood and adolescence (Pollak, 1976). De la Croix (1996) builds on this tradition by specifying interdependent preferences where parents’ consumption levels serve as the reference point for their children’s future consumption. Our framework aligns with Pollak’s work by directly setting the reference benchmark to the birth year. Our focus is on examining the properties of the class of functions

that, grounded in these cohort-specific behaviors, allow researchers to transcend the functional restrictions of benchmark-independent preferences when seeking consistency with balanced growth paths.

3. The Current Paradigm: The Restriction of Benchmark-Independent Preferences

In this section, we define the balanced growth path (BGP) with trendless hours and discuss the well-known preferences consistent with a BGP within the domain of felicity functions where consumption and leisure are evaluated in absolute terms, absent aggregate benchmarks. These established results by King et al. (1988a,b), along with the extension by Boppart and Krusell (2020) to functional forms compatible with secular declines in hours worked (discussed in Section 4.5), represent the current paradigm for balanced growth preferences: a framework with benchmark-independent felicity functions.

We focus on environments where, along a BGP, the time path of an economic variable x is given by $x_t = (1 + \gamma_x)^t x_0$, where γ_x is its constant growth rate and x_0 is its initial value. The textbook definition of a BGP incorporating the trendless behavior of hours is as follows:

Balanced Growth Path (King et al. (1988a,b)). *A balanced growth path (BGP) is a growth path with a constant interest rate (r) along which:*

- (a) *Consumption (c) and labor income (wh) per capita grow at the same constant rate ($\gamma_c = \gamma_{wh} = \gamma$),*

$$c_t = (1 + \gamma)^t c_0 \quad \text{and} \quad w_t h_t = (1 + \gamma)^t w_0 h_0 \quad (1)$$

- (b) *Hours per capita remain constant ($\gamma_h = 0$)*

$$h_t = h_0 \quad (2)$$

Remark 1. *The combination of (a) and (b) implies that $\gamma_c = \gamma_w = \gamma$ along the BGP.*

3.1. Infinite-Horizon Models with benchmark-independent Felicity Functions

Consider a representative agent that chooses consumption c and labor supply h (or leisure $\ell = 1 - h$) to solve an infinite horizon model with a technologically determined growth rate. Let the benchmark-independent felicity function of the representative agent, $u(c, h)$, be twice continuously differentiable, strictly increasing in consumption, and strictly concave in both arguments. These are regular assumptions ensuring positive and diminishing marginal utility. As established by King et al. (1988a,b), achieving a BGP in this setting requires specific functional restrictions.

Specifically, the felicity function—without aggregate benchmarks—is consistent with BGP if and only if it takes the form $u(c_t, h_t) = \frac{c_t^{1-\sigma}}{1-\sigma} v(h_t)$ if $\sigma \neq 1$ or $u(c_t, h_t) = \ln c_t + v(h_t)$ if $\sigma = 1$.

The infinite-horizon model serves as a convenient approximation for overlapping-generations (OLG) frameworks; consequently, the classic restrictions on benchmark-independent preferences established by King et al. (1988a) carry over directly to the OLG case. The following section formally defines our OLG environment. We operate within this OLG framework throughout the main text, as it is a natural setting for our advocated preferences—specifically, birth-level normalization—and we provide analogous infinite-horizon results in the Appendix.

3.2. OLG Model

Consider an economy with overlapping generations and finitely lived agents with ages ranging from $i \in \{0, \dots, I\}$. Let $t \in \{0, \dots, \infty\}$ denote discrete calendar time. An agent of age i in period t belongs to the cohort born at time $\tau = t - i$. We assume the evolution of the population is deterministic and that the age composition remains constant over time. Under these conditions, a finitely lived agent born in period τ has a benchmark-independent felicity function with the standard regularity properties and solves:

$$\max_{\{c_t^i, h_t^i\}} \sum_{i=0}^I \beta^i u(c_t^i, h_t^i) \quad (3)$$

$$\text{s.t.} \quad \sum_{i=0}^I \frac{c_t^i}{(1+r)^{t-\tau}} = \sum_{i=0}^I \frac{w_t h_t^i}{(1+r)^{t-\tau}} + a_0 \quad (4)$$

with a technologically determined growth rate.²

In this OLG context, the definition of BGP additionally captures identical life-cycle behavior of consumption and labor income across cohorts. Specifically, the fact that per-capita consumption and labor income grow at the same constant rate along the BGP as in (1) implies that for all ages i :

$$\frac{c_{t+1}^i}{c_t^i} = 1 + \gamma \quad \text{and} \quad \frac{w_{t+1} h_{t+1}^i}{w_t h_t^i} = 1 + \gamma \quad (5)$$

or equivalently $c_t^i = (1 + \gamma)^t c_0^i$ and $w_t h_t^i = (1 + \gamma)^t w_0 h_0^i$. And the fact that aggregate hours per capita remain constant along the BGP as in (2) implies that for all ages i :

$$\frac{h_{t+1}^i}{h_t^i} = 1 \quad (6)$$

²Wages are identical across age groups—a feature that arises with perfect substitutability of labor across age groups in the production function. This is innocuous for our purposes of determining the properties of the felicity function that are consistent with a BGP. Our results carry over when assuming an efficiency-wage age profile or age-biased technology.

or equivalently $h_t^i = h_0^i$. This aggregate behavior described in (5) and (6) in turn implies that the life-cycle profiles of consumption (and labor income) are the same across all cohorts τ along the BGP. That is, for all ages i

$$\frac{c_{t+1}^{i+1}}{c_t^i} = (1 + \gamma)(1 + g_{i,c}) \quad \text{and} \quad \frac{h_{t+1}^{i+1}}{h_t^i} = 1 + g_{i,h} \quad (7)$$

or equivalently $c_t^i = (1 + \gamma)^i \Pi_{t=0}^{i-1} (1 + g_{i,c}) c_{t-i}^0$ and $h_t^i = \Pi_{t=0}^{i-1} (1 + g_{i,h}) h_{t-i}^0$ with $g_{i,c}$ and $g_{i,h}$ being, respectively, the life-cycle consumption and labor income growth between the ages i and $i + 1$ net of aggregate growth γ . It is straightforward to show that the properties of the benchmark-independent felicity function consistent with a BGP in the OLG model are identical to those in the infinite-horizon model.

We now turn to our core contribution. We introduce a broader class of functions consistent with balanced growth without necessitating any specific functional form, effectively decoupling the shape of the felicity function from long-run growth requirements. In Section 4, we show that this class nests the specification in King et al. (1988a,b). Furthermore, in an extension in Section 4.5, we show that it also nests Boppart and Krusell (2020) as a special case.

4. A General Class of BGP-Consistent Preferences with Aggregate Benchmarks

We propose a broad class of preferences defined over individual consumption c , labor supply h , and aggregate consumption C . We show that the property of *pairwise homogeneity (PH)* in the pair (c, C) is both necessary and sufficient for balanced growth. All preferences satisfying this homogeneity property are consistent with a BGP without further restrictions on the functional form of the felicity function beyond standard regularity conditions. In this manner, our approach decouples the shape of the utility function from the requirements of a BGP. We demonstrate that KPR preferences constitute a sub-manifold of PH—a knife-edge case that breaks the decoupling. We extend this argument to *independent pairwise homogeneity* in Section 4.5 to account for the secular decline in labor supply (Boppart and Krusell, 2020).

4.1. Pairwise Homogeneous Preferences

Let $c_t \in \mathbb{R}^{++}$ denote individual consumption at time t , and let $h_t \in (0, 1]$ be the individual labor supply, with leisure defined as $\ell_t = 1 - h_t$. Let $C_\tau \in \mathbb{R}^{++}$ (for $\tau \leq t$) represent the per-capita consumption benchmark, which grows at the same rate as c_t along the balanced growth path (BGP). We consider the following general class of preferences:

$$U(c_t, h_t; C_\tau) \quad (8)$$

These preferences are assumed to be twice continuously differentiable in all arguments and satisfy standard properties, including strictly positive and diminishing marginal utility with respect to individual choices c_t and h_t . Furthermore, the conditions $U_C < 0$, $U_{CC} > 0$, and $U_{cC} > 0$ characterize the benchmark C_τ as a relative reference point for the standard of living. In particular, the positive cross-derivative U_{cC} ensures that the benchmark increases the marginal utility of individual consumption.

Definition 1 (*PH(k) Preferences*). *The preferences $U(c, h; C)$ are pairwise homogeneous of degree k in the pair (c, C) if, for any scaling factor $\lambda > 0$:*

$$U(\lambda c, h; \lambda C) = \lambda^k U(c, h; C) \quad (9)$$

This implies that preferences exhibit homogeneity of degree k specifically with respect to individual consumption c and its aggregate counterpart C , while h remains unscaled. Setting the scaling factor to $\lambda = C_\tau^{-1}$, we obtain:

$$U(c_t, h_t; C_\tau) = C_\tau^k u\left(\frac{c_t}{C_\tau}, h_t\right) \quad (10)$$

where $u(\cdot, \cdot)$ is the felicity function capturing the intensive formulation of preferences, that is, $u\left(\frac{c_t}{C_\tau}, h_t\right) \equiv U\left(\frac{c_t}{C_\tau}, h_t; 1\right)$. Thus, this felicity function follows the properties of U and satisfies the following regularity conditions:

$$u \in C^2, \quad u : \mathbb{R}_{++} \times [0, 1] \rightarrow \mathbb{R}, \quad u_c > 0, \quad u_{cc} < 0, \quad u_h < 0, \quad u_{hh} > 0 \quad (11)$$

The degree of homogeneity, k , in (10) governs how the level of utility scales with the overall size of the economy. If $k = 0$, the utility is scale-invariant, meaning absolute well-being depends only on relative consumption c/C . If $k \neq 0$, the absolute level of utility shifts as the economy grows, though the underlying preference ordering over (c, h) remains unchanged for a given C_τ . This felicity function indicates that a reference level of consumption—such as the level at birth if $\tau = t - i$ —serves as the unit of account. An agent’s relative well-being remains invariant to the absolute scale of the economy, provided consumption maintains a constant proportion to the norms that prevailed at that reference time.

Lemma 1. *The preferences $U(c, h; C)$ are PH(k) in (c, C) if and only if the marginal utility $U_c(c, h; C)$ is PH(k-1) and the marginal utility $U_h(c, h; C)$ is PH(k) in the pair (c, C) .*

Proof. We prove sufficiency and necessity separately:

- (a) The sufficiency (“if”) direction follows from Euler’s Theorem for homogeneous functions. If the partial derivatives with respect to the scaling variables c and C are PH($k - 1$), and the

derivative with respect to h is $PH(k)$, then the total differential $dU = U_c dc + U_h dh$ must scale such that the primitive function U is $PH(k)$.³

- (b) The necessity (“only if”) direction follows from differentiating the scaling identity $U(\lambda c, h, \lambda C) = \lambda^k U(c, h, C)$ with respect to c and h . Using the intensive formulation (10), we obtain:

$$U_c(c_t, h_t; C_\tau) = C_\tau^{k-1} u_c \left(\frac{c_t}{C_\tau}, h_t \right), \quad (12)$$

$$U_h(c_t, h_t; C_\tau) = C_\tau^k u_h \left(\frac{c_t}{C_\tau}, h_t \right), \quad (13)$$

where u_c and u_h denote the partial derivatives of $u(\cdot, \cdot)$ with respect to its first and second arguments. Since these arguments depend only on the ratio c_t/C_τ , they are homogeneous of degree zero in (c, C) . Thus, U_c and U_h are $PH(k-1)$ and $PH(k)$, respectively.

□

Lemma 2. *The ratio of $U_c(c, h; C)$ evaluated at two arbitrary allocations (c_t, h_t) and (c_{t+1}, h_{t+1}) satisfying $(c_{t+1}, h_{t+1}) = ((1 + \gamma)c_t, h_t)$ and sharing a common benchmark C_τ is pairwise homogeneous of degree zero, $PH(0)$, in the pair (c, C) —meaning it remains scale-invariant under simultaneous scaling by a factor $\lambda > 0$ in c and C :*

$$RMU = \frac{U_c(\lambda c_t, h_t; \lambda C_\tau)}{U_c(\lambda c_{t+1}, h_{t+1}; \lambda C_\tau)} = \frac{U_c(\lambda c_t, h_t; \lambda C_\tau)}{U_c(\lambda(1 + \gamma)c_t, h_t; \lambda C_\tau)} = \frac{U_c(c_t, h_t; C_\tau)}{U_c((1 + \gamma)c_t, h_t; C_\tau)}$$

if and only if the marginal utility of consumption, $U_c(c, h; C)$, is pairwise homogeneous of degree $k-1$ in the pair (c, C) ; that is, $U_c(\lambda c, h; \lambda C) = \lambda^{k-1} U_c(c, h; C)$ for some constant $k-1 \in \mathbb{R}$.

Proof. We prove sufficiency and necessity separately:

- (a) U_c is $PH(k-1)$ in $(c, C) \implies$ RMU is $PH(0)$ in (c, C) (sufficiency): This (the “if” direction) follows immediately from substituting a homogeneous function of degree $k-1$ into the ratio, where the scaling factor λ^{k-1} cancels out from both the numerator and denominator.

³Specifically, if U_c and U_h are $PH(k-1)$, then for any $\lambda > 0$:

$$\begin{aligned} kU(\lambda c, h, \lambda C) &= (\lambda c)U_c(\lambda c, h, \lambda C) + (\lambda C)U_h(\lambda c, h, \lambda C) \\ &= \lambda c[\lambda^{k-1} U_c(c, h, C)] + \lambda C[\lambda^{k-1} U_h(c, h, C)] \\ &= \lambda^k [cU_c(c, h, C) + CU_h(c, h, C)] = \lambda^k [kU(c, h, C)]. \end{aligned}$$

Dividing by k yields the $PH(k)$ scaling property for U . Since U_h is $PH(k)$, this scaling property holds for any h , confirming that U is $PH(k)$ in the pair (c, C) .

- (b) RMU is PH(0) in $(c, C) \implies U_c$ is PH($k - 1$) in (c, C) (necessity): To prove the “only if” direction, we assume that the ratio is scale-invariant, i.e., PH(0) under the scaling factor $\lambda > 0$. Equating the baseline ratio with its scaled counterpart yields:

$$\frac{U_c(c_t, h_t; C_\tau)}{U_c((1 + \gamma)c_t, h_t; C_\tau)} = \frac{U_c(\lambda c_t, h_t; \lambda C_\tau)}{U_c(\lambda(1 + \gamma)c_t, h_t; \lambda C_\tau)}$$

This identity shows that shifting the individual allocation argument leaves the proportional scaling response of the marginal utility unchanged. Consequently, this scaling multiplier is independent of individual allocations and must depend strictly on λ . Defining this function as $\phi(\lambda)$, we can write

$$U_c(\lambda c_t, h_t; \lambda C_\tau) = \phi(\lambda) U_c(c_t, h_t; C_\tau)$$

Sequential scaling by arbitrary factors $\lambda_1, \lambda_2 > 0$ yields Cauchy’s multiplicative functional equation $\phi(\lambda_1 \lambda_2) = \phi(\lambda_1) \phi(\lambda_2)$. Given that $U \in C^2$, $\phi(\lambda)$ is continuous. Its unique non-trivial continuous solution is a power function $\phi(\lambda) = \lambda^{k-1}$ for some constant $k - 1 \in \mathbb{R}$ (see Sydsæter et al. (2008)). Substituting this back provides $U_c(\lambda c_t, h_t; \lambda C_\tau) = \lambda^{k-1} U_c(c_t, h_t; C_\tau)$, completing the proof that U_c must be pairwise homogeneous of degree $k - 1$ in (c, C) .

□

4.2. Balanced Growth with PH(k) Preferences

In what follows, we focus on an OLG economy identical to the one described in the previous section, and subsequently extend the analysis to the infinite horizon in Appendix A.1. Taking the scaling factor C_τ as given, an agent of age i at time t who enjoys a utility $U(c_t^i, h_t^i; C_\tau)$ chooses individual consumption c_t^i and labor supply h_t^i by solving:

$$\max_{\{c_t^i, h_t^i\}_{i=0}^I} \sum_{i=0}^I \beta^i U(c_t^i, h_t^i; C_\tau) \tag{14}$$

$$\text{s.t.} \quad \sum_{i=0}^I \frac{c_{t+i}^i}{(1+r)^i} = \sum_{i=0}^I \frac{w_{t+i} h_{t+i}^i}{(1+r)^i} + a_0 \tag{15}$$

given a_0 and a technologically determined growth rate. The intratemporal first-order condition for an agent of age i in period t is $U_h(c_t^i, h_t^i; C_\tau) = w_t U_c(c_t^i, h_t^i; C_\tau)$. The intertemporal first-order condition for an agent transitioning from age i at time t to age $i + 1$ at $t + 1$ is $U_c(c_t^i, h_t^i; C_\tau) = \beta(1+r) U_c(c_{t+1}^{i+1}, h_{t+1}^{i+1}; C_\tau)$.

Proposition 1. *The OLG economy described by (14)–(15) with preferences $U(c_t, h_t; C_\tau)$ admits a BGP with trendless hours (5)–(7) if and only if the preferences are $PH(k)$ in (c, C) ; that is,*

$$U(c_t, h_t; C_\tau) = C_\tau^k u\left(\frac{c_t}{C_\tau}, h_t\right) \quad (16)$$

for any felicity function $u(\cdot, \cdot)$ satisfying the regularity conditions in (11).

Proof. (a) $PH(k) \Rightarrow$ BGP. The proof of the sufficiency of $PH(k)$ proceeds in four steps:

1. *Intertemporal condition:* Plugging the $PH(k)$ preferences (16) into the objective function (14), we obtain the intertemporal condition:

$$C_\tau^k u_{\bar{c}}\left(\frac{c_t^i}{C_\tau}, h_t^i\right) \frac{1}{C_\tau} = \beta (1+r) C_\tau^k u_{\bar{c}}\left(\frac{c_{t+1}^{i+1}}{C_\tau}, h_{t+1}^{i+1}\right) \frac{1}{C_\tau}$$

where the marginal utility of consumption is $u_c\left(\frac{c_t^i}{C_\tau}, h_t^i\right) = u_{\bar{c}}\left(\frac{c_t^i}{C_\tau}, h_t^i\right) \frac{1}{C_\tau}$. Equating the intertemporal MRS for an agent of age i in period t (cohort τ) to that of an agent of the same age i one period forward (cohort $\tau + 1$):

$$\underbrace{\frac{C_\tau^k u_{\bar{c}}\left(\frac{c_t^i}{C_\tau}, h_t^i\right) \frac{1}{C_\tau}}{\beta C_\tau^k u_{\bar{c}}\left(\frac{c_{t+1}^{i+1}}{C_\tau}, h_{t+1}^{i+1}\right) \frac{1}{C_\tau}}}_{\text{Intertemporal } MRS(c_t^i, c_{t+1}^{i+1})} = \underbrace{\frac{C_{\tau+1}^k u_{\bar{c}}\left(\frac{c_{t+1}^i}{C_{\tau+1}}, h_{t+1}^i\right) \frac{1}{C_{\tau+1}}}{\beta C_{\tau+1}^k u_{\bar{c}}\left(\frac{c_{t+2}^{i+1}}{C_{\tau+1}}, h_{t+2}^{i+1}\right) \frac{1}{C_{\tau+1}}}}_{\text{Intertemporal } MRS(c_{t+1}^i, c_{t+2}^{i+1})} = 1 + r$$

which after cancelling scaling factors becomes:

$$\underbrace{\frac{u_{\bar{c}}\left(\frac{c_t^i}{C_\tau}, h_t^i\right)}{\beta u_{\bar{c}}\left(\frac{c_{t+1}^{i+1}}{C_\tau}, h_{t+1}^{i+1}\right)}}_{\text{Intertemporal } MRS(c_t^i, c_{t+1}^{i+1})} = \underbrace{\frac{u_{\bar{c}}\left(\frac{c_{t+1}^i}{C_{\tau+1}}, h_{t+1}^i\right)}{\beta u_{\bar{c}}\left(\frac{c_{t+2}^{i+1}}{C_{\tau+1}}, h_{t+2}^{i+1}\right)}}_{\text{Intertemporal } MRS(c_{t+1}^i, c_{t+2}^{i+1})} = 1 + r \quad (17)$$

Our goal is to show the stationarity of the intertemporal condition along the BGP. With r constant, this is equivalent to demonstrating that the intertemporal MRS throughout the life cycle is invariant across cohorts.

2. *Stationarity of the intertemporal condition:* Plug the BGP constraints (5)–(7) into (17)

to substitute in c_t^i , h_t^i and C_τ :

$$\frac{u_{\bar{c}} \left(\frac{c_t^i}{C_\tau}, h_t^i \right)}{\beta u_{\bar{c}} \left(\frac{(1+g_{i,c})(1+\gamma)c_t^i}{C_\tau}, (1+g_{i,h})h_t^i \right)} = \frac{u_{\bar{c}} \left(\frac{(1+\gamma)c_t^i}{(1+\gamma)C_\tau}, h_t^i \right)}{\beta u_{\bar{c}} \left(\frac{(1+g_{i,c})(1+\gamma)^2 c_t^i}{(1+\gamma)C_\tau}, (1+g_{i,h})h_t^i \right)}$$

Canceling the growth terms $(1+\gamma)$ within arguments in the marginal utilities of the $\tau+1$ cohort we obtain:

$$\underbrace{\frac{u_{\bar{c}} \left(\frac{c_t^i}{C_\tau}, h_t^i \right)}{\beta u_{\bar{c}} \left(\frac{(1+g_{i,c})(1+\gamma)c_t^i}{C_\tau}, (1+g_{i,h})h_t^i \right)}}_{\text{Intertemporal } MRS(c_t^i, c_{t+1}^i)} = \underbrace{\frac{u_{\bar{c}} \left(\frac{c_t^i}{C_\tau}, h_t^i \right)}{\beta u_{\bar{c}} \left(\frac{(1+g_{i,c})(1+\gamma)c_t^i}{C_\tau}, (1+g_{i,h})h_t^i \right)}}_{\text{Intertemporal } MRS(c_{t+1}^i, c_{t+2}^i)} = \text{Constant},$$

which is satisfied for all cohorts τ . Consequently, the intertemporal condition remains stationary under any regular $PH(k)$ function.

3. *Intratemporal condition:* Plugging the $PH(k)$ preferences (16) into the objective function (14), we obtain the intratemporal condition:

$$C_\tau^k u_h \left(\frac{c_t^i}{C_\tau}, h_t^i \right) = w_t C_\tau^k u_{\bar{c}} \left(\frac{c_t^i}{C_\tau}, h_t^i \right) \frac{1}{C_\tau} \quad (18)$$

Dividing the intratemporal MRS for an agent of age i in period t by that of an agent of the same age i in the subsequent cohort and rearranging we obtain:

$$\frac{C_\tau^k u_{\bar{c}} \left(\frac{c_t^i}{C_\tau}, h_t^i \right) \frac{1}{C_\tau}}{C_\tau^k u_h \left(\frac{c_t^i}{C_\tau}, h_t^i \right)} = \frac{w_{t+1}}{w_t} \frac{C_{\tau+1}^k u_{\bar{c}} \left(\frac{c_{t+1}^i}{C_{\tau+1}}, h_{t+1}^i \right) \frac{1}{C_{\tau+1}}}{C_{\tau+1}^k u_h \left(\frac{c_{t+1}^i}{C_{\tau+1}}, h_{t+1}^i \right)}$$

Intratemporal $MRS(c_t^i, h_t^i)$
Intratemporal $MRS(c_{t+1}^i, h_{t+1}^i)$

which after canceling scaling factors becomes:

$$\frac{\frac{u_{\bar{c}}\left(\frac{c_t^i}{C_\tau}, h_t^i\right) \frac{1}{C_\tau}}{u_h\left(\frac{c_t^i}{C_\tau}, h_t^i\right)}}{\underbrace{\hspace{1.5cm}}_{\text{Intratemporal } MRS(c_t^i, h_t^i)}} = \frac{w_{t+1}}{w_t} \frac{\frac{u_{\bar{c}}\left(\frac{c_{t+1}^i}{C_{\tau+1}}, h_{t+1}^i\right) \frac{1}{C_{\tau+1}}}{u_h\left(\frac{c_{t+1}^i}{C_{\tau+1}}, h_{t+1}^i\right)}}{\underbrace{\hspace{1.5cm}}_{\text{Intratemporal } MRS(c_{t+1}^i, h_{t+1}^i)}} \quad (19)$$

We are after the stationarity of the intratemporal condition along the BGP. Since wages grow, the lifecycle intratemporal MRS must exactly offset this wage growth to ensure invariance across cohorts.

4. *Stationarity of the intratemporal condition:* Rearrange scaling factors outside the marginal utilities and plug the BGP constraints (5)–(7) into (19) to substitute in c_t^i , h_t^i , C_τ and w_t to obtain:

$$\frac{u_{\bar{c}}\left(\frac{c_t^i}{C_\tau}, h_t^i\right)}{u_h\left(\frac{c_t^i}{C_\tau}, h_t^i\right)} = \frac{(1+\gamma)w_t}{w_t} \frac{C_\tau}{(1+\gamma)C_\tau} \frac{u_{\bar{c}}\left(\frac{(1+\gamma)c_t^i}{(1+\gamma)C_\tau}, h_t^i\right)}{u_h\left(\frac{(1+\gamma)c_t^i}{(1+\gamma)C_\tau}, h_t^i\right)}$$

Canceling w_t , scaling factors C_τ and the growth terms $(1+\gamma)$ yields:

$$\frac{\frac{u_{\bar{c}}\left(\frac{c_t^i}{C_\tau}, h_t^i\right)}{u_h\left(\frac{c_t^i}{C_\tau}, h_t^i\right)}}{\underbrace{\hspace{1.5cm}}_{\text{Intratemporal } MRS(c_t^i, h_t^i)}} = \frac{\frac{u_{\bar{c}}\left(\frac{c_t^i}{C_\tau}, h_t^i\right)}{u_h\left(\frac{c_t^i}{C_\tau}, h_t^i\right)}}{\underbrace{\hspace{1.5cm}}_{\text{Intratemporal } MRS(c_{t+1}^i, h_{t+1}^i)}} = \text{Constant}$$

This equality holds for all cohorts τ . Consequently, the intratemporal condition is stationary across the entire sequence of generations for any regular $PH(k)$ function for any regular felicity function u .

Thus, if U is $PH(k)$ in (c, C) , then both the intertemporal condition and the intratemporal condition are stationary across all cohorts, that is, the growing economy admits a BGP.

- (b) $\text{BGP} \Rightarrow PH(k)$. The proof of the necessity of $PH(k)$ consists of two needs:

1. *Intertemporal MRS must be stationary:* The intertemporal condition for two consecu-

tive cohorts is:

$$\frac{U_c(c_t^i, h_t^i; C_\tau)}{\beta U_c(c_{t+1}^{i+1}, h_{t+1}^{i+1}; C_\tau)} = \frac{U_c(c_{t+1}^i, h_{t+1}^i; C_{\tau+1})}{\beta U_c(c_{t+2}^{i+1}, h_{t+2}^{i+1}; C_{\tau+1})} = 1 + r$$

which, along the BGP, becomes:

$$\begin{aligned} \frac{U_c(c_t^i, h_t^i; C_\tau)}{\beta U_c((1+\gamma)(1+g_{i,c})c_t^i, (1+g_{i,h})h_t^i; C_\tau)} &= \frac{U_c((1+\gamma)c_t^i, h_t^i; (1+\gamma)C_\tau)}{\beta U_c((1+\gamma)^2(1+g_{i,c})c_t^i, (1+g_{i,h})h_t^i; (1+\gamma)C_\tau)} \\ &= \frac{U_c((1+\gamma)c_t^i, h_t^i; (1+\gamma)C_\tau)}{\beta U_c((1+\gamma)(1+\gamma)(1+g_{i,c})c_t^i, (1+g_{i,h})h_t^i; (1+\gamma)C_\tau)} = 1 + r \end{aligned} \quad (20)$$

This equality implies that along the BGP, the intertemporal MRS must be PH(0) in the pair (c, C) —that is, scale-invariant with respect to the growing variables (c, C) —requiring the growth terms $(1+\gamma)$ to cancel out uniformly. Consequently, the scaling multiplier is independent of individual allocations and depends strictly on the aggregate scaling factor $(1+\gamma)$, implying $U_c((1+\gamma)c_t^i, h_t^i; (1+\gamma)C_\tau) = \phi(1+\gamma)U_c(c_t^i, h_t^i; C_\tau)$. Hence, evaluating the intertemporal MRS of cohort $\tau + 1$ yields:

$$\begin{aligned} \frac{U_c((1+\gamma)c_t^i, h_t^i; (1+\gamma)C_\tau)}{\beta U_c((1+\gamma)(1+\gamma)(1+g_{i,c})c_t^i, (1+g_{i,h})h_t^i; (1+\gamma)C_\tau)} &= \frac{\phi(1+\gamma)U_c(c_t^i, h_t^i; C_\tau)}{\beta \phi(1+\gamma)U_c((1+\gamma)(1+g_{i,c})c_t^i, (1+g_{i,h})h_t^i; C_\tau)} \\ &= \frac{U_c(c_t^i, h_t^i; C_\tau)}{\beta U_c((1+\gamma)(1+g_{i,c})c_t^i, (1+g_{i,h})h_t^i; C_\tau)} = 1 + r \end{aligned} \quad (21)$$

which guarantees the stationarity of the intertemporal MRS for all cohorts. Since the discount factor β is a constant, the stationarity of the intertemporal MRS is equivalent to the stationarity of the ratio of marginal utilities across successive periods. Thus, by Lemma 2 (part (b)), this structural configuration necessitates that U_c is PH($k-1$) in (c, C) so that the aggregate scaling factors mapping cohort τ to $\tau + 1$ factor out uniformly from the marginal utilities, establishing that $\phi(1+\gamma) = (1+\gamma)^{k-1}$.

Thus, the stationarity of the intertemporal MRS implies that U_c is required to be PH($k-1$) in (c, C) .

2. *Intratemporal MRS must offset the grow rate of wages:* The intratemporal condition

for two consecutive cohorts is:

$$\frac{U_c(c_t^i, h_t^i; C_\tau)}{U_h(c_t^i, h_t^i; C_\tau)} = \frac{w_{t+1}}{w_t} \frac{U_c(c_{t+1}^i, h_{t+1}^i; C_{\tau+1})}{U_h(c_{t+1}^i, h_{t+1}^i; C_{\tau+1})}$$

which along the BGP is,

$$\frac{U_c(c_t^i, h_t^i; C_\tau)}{U_h(c_t^i, h_t^i; C_\tau)} = (1 + \gamma) \frac{U_c((1 + \gamma)c_t^i, h_t^i; (1 + \gamma)C_\tau)}{U_h((1 + \gamma)c_t^i, h_t^i; (1 + \gamma)C_\tau)} \quad (22)$$

Given that wages grow at rate $(1 + \gamma)$ and that U_c is $PH(k - 1)$ in (c, C) (from Step 1), stationarity of the intratemporal condition (22) requires U_h to be homogeneous of degree k in (c, C) to offset the growth in wages and marginal utility:

$$\frac{U_c(c_t^i, h_t^i; C_\tau)}{U_h(c_t^i, h_t^i; C_\tau)} = (1 + \gamma) \frac{(1 + \gamma)^{k-1}}{(1 + \gamma)^k} \frac{U_c(c_t^i, h_t^i; C_\tau)}{U_h(c_t^i, h_t^i; C_\tau)} = \frac{U_c(c_t^i, h_t^i; C_\tau)}{U_h(c_t^i, h_t^i; C_\tau)} \quad (23)$$

Finally, collecting the joint requirements that U_c needs to be $PH(k - 1)$ in (c, C) and U_h needs to be $PH(k)$ in (c, C) implies that the utility function U needs to be $PH(k)$ in (c, C) , as established in Lemma 1. □

Proposition 1 demonstrates that BGP consistency is driven entirely by the property of *pairwise homogeneity* in (c, C) , rather than by any specific functional form. Under $PH(k)$, any felicity function satisfying standard regularity conditions is consistent with a BGP.

4.3. Extension to Infinite-Horizon

In Appendix A.1, we provide the infinite-horizon counterpart to the OLG proposition. This model represents the limiting case of our economy; while the first-order conditions for balanced growth remain identical across both frameworks, the infinite-horizon limit necessitates an additional transversality condition to ensure a well-defined equilibrium. We show that $PH(k)$ remains the fundamental property for preferences required for balanced growth.

4.4. The Current Paradigm: A Sub-Manifold of $PH(k)$

We argue that the “functional-form discipline” of the current paradigm does not stem from an inherent property of balanced growth but, instead, arises from the introduction of an independent modeling choice not necessary for a BGP: time invariance in the utility function.

To see this, consider the intersection of two distinct properties: *pairwise homogeneity* of $U(c, h; C)$ in (c, C) (necessary for a BGP as per Proposition 1) and *benchmark-independence*

$U(c, h; C) = U(c, h)$. That is, impose benchmark-independence on the more general class of functions that satisfy *PH*:

$$\underbrace{U(c_t, h_t; C_\tau) = C_\tau^k u\left(\frac{c_t}{C_\tau}, h_t\right)}_{PH(k) \text{ Preferences}} \underbrace{=}_{\text{benchmark-independence Constraint}} U(c_t, h_t) \quad (24)$$

which implies that the marginal utility of consumption $U_c(c, h; C)$ satisfies:

$$C^{k-1} u_c\left(\frac{c}{C}, h\right) = U_c(c, h) \quad (25)$$

Because the right-hand side of (25) does not depend on C , its partial derivative with respect to C must be zero. Consequently, the partial derivative of the left-hand side must also be zero. Applying the product rule and chain rule, and letting $z \equiv c/C$:

$$(k-1)C^{k-2}u_z(z, h) + C^{k-1}u_{zz}(z, h)\left(-\frac{c}{C^2}\right) = 0 \quad (26)$$

Dividing by C^{k-2} yields $(k-1)u_z(z, h) - zu_{zz}(z, h) = 0$, which implies:

$$\frac{zu_{zz}(z, h)}{u_z(z, h)} = k-1 \quad (27)$$

The differential equation in (27) implies that $u_z(z, h)$ is isoelastic in z , which recovers the result of the current paradigm (King et al., 1988a,b). Specifically, integrating both sides with respect to z yields:

$$u_z(z, h) = z^{k-1}v(h) \quad (28)$$

where $v(h)$ is a constant of integration. Substituting this back into (25) and using the definition $z = c/C$, we find:

$$U_c(c, h; C) = C^{k-1} \underbrace{\left(\frac{c}{C}\right)^{k-1} v(h)}_{=u_z(z, h) \text{ from Eq. (28)}} = c^{k-1}v(h) = U_c(c, h) \quad (29)$$

This confirms that benchmark-independence restricts the marginal utility to a power function, effectively imposing the standard CRRA (Constant Relative Risk Aversion) structure required by (King et al., 1988a,b). Importantly, (27) reveals that imposing benchmark-independence requires the elasticity not only to be constant, but exactly equal to $k-1$. The reason is that only under this specific shape do the aggregate benchmark terms (C_τ) in the marginal utility of our general formulation, $U_c(c, h; C)$, mathematically cancel out (see equation (29)), rendering the specification equivalent to a benchmark-independent one. In other words, to remove C_τ from $U_c(c, h; C)$ —

achieving the benchmark-independent counterpart $U_c(c, h)$ —the current paradigm forces the burden of BGP consistency entirely onto the curvature of the individual felicity function. It is only this unique configuration that the functional form effectively “swallows” the aggregate scaling factor. This creates the illusion that BGP consistency is a property of the felicity function’s shape, rather than the underlying homogeneity of the environment.⁴

Summing up, in the general $PH(k)$ class, BGP consistency is a structural property of the *pairwise* relationship between individual and aggregate consumption: the aggregate reference C_τ performs the “lifting” required to keep the intertemporal *MRS* stationary while simultaneously balancing the growing real wage. This allows any regular felicity function to be consistent with balanced growth. In contrast, by imposing the modeling choice of benchmark-independence, the current paradigm forces the curvature of u to fulfill a role naturally handled by the aggregate benchmark, breaking the *decoupling* of growth consistency from the shape of preferences. This rigid dependency limits the set of admissible preferences to those in King et al. (1988a,b)—a knife-edge sub-manifold of the $PH(k)$ class.

4.5. Extension to Trends in Hours Worked

Recent empirical research has highlighted a persistent decline in hours worked across time and space (Bick et al., 2018; Boppart and Krusell, 2020). In particular, Boppart and Krusell (2020) uncover specific functional forms of benchmark-independent felicity functions that are consistent with a BGP accompanied by a long-run decline in hours worked ($\gamma_h < 0$). To accommodate trends in hours, we introduce a broader class of preferences with *independent pairwise homogeneity* in both pairs (c, C) and (h, H) , which nests the single pairwise homogeneity of the previous section.

The labor benchmark H_τ has a direct behavioral reading: it is the *work norm* of a cohort’s formative environment—how much, and by whom, it is normal to work. Just as C_τ anchors the consumption standard, H_τ anchors the labor standard. It is this benchmark that performs the leisure-side “lifting”: by absorbing the secular growth of the real wage, H_τ frees the Frisch elasticity from the curvature of u and lets hours follow any trend. The evidence in Section 6.3 is about precisely this object—the cohort-and-place specific norm of (female) work.

The following definition of the BGP accommodates aggregate labor trends.

Balanced Growth Path (Boppart and Krusell (2020)). *A balanced growth path is a growth path with a constant interest rate (r) along which:*

⁴An alternative proof of necessity within the benchmark-independent setting based on Cauchy functional equations is provided in Appendix B to draw a parallel with the proof strategy for Proposition 1. Specifically, we contrast the necessity step under pairwise homogeneity (part (b.1) of Proposition 1)—where growth scaling is absorbed by the aggregate benchmark—with the benchmark-independent paradigm, which forces the individual marginal utility to absorb this structural weight via Cauchy’s functional equation.

(a) Consumption (c) and labor income (wh) per capita grow at the same constant rate ($\gamma_c = \gamma_{wh} = \gamma$),

$$c_t = (1 + \gamma)^t c_0 \quad \text{and} \quad w_t h_t = (1 + \gamma)^t w_0 h_0 \quad (30)$$

(b) Hours per capita (h) decline at a constant rate ($\gamma_h < 0$),

$$h_t = (1 + \gamma_h)^t h_0 \quad (31)$$

Remark 2. The combination of (a) and (b) implies that $\gamma_w > \gamma_c = \gamma$ along this path.

From an OLG perspective, this definition ensures that life-cycle profiles of consumption and labor income remain stable across successive cohorts. Specifically, for all ages i :

$$\frac{c_{t+1}^i}{c_t^i} = 1 + \gamma \quad \text{and} \quad \frac{w_{t+1} h_{t+1}^i}{w_t h_t^i} = 1 + \gamma \quad (32)$$

and

$$\frac{h_{t+1}^i}{h_t^i} = 1 + \gamma_h \quad (33)$$

with life-cycle profiles satisfying, for all ages i :

$$\frac{c_{t+1}^{i+1}}{c_t^i} = (1 + \gamma)(1 + g_{i,c}) \quad \text{and} \quad \frac{w_{t+1} h_{t+1}^{i+1}}{w_t h_t^i} = (1 + \gamma_w)(1 + \gamma_h)(1 + g_{i,h}) \quad (34)$$

4.5.1. Independent Pairwise Homogeneous Preferences

In addition to C_τ , let H_τ represent the aggregate per-capita labor supply, which grows at the same rate as h_t along the BGP, with $\tau \leq t$. We consider a general class of preferences:

$$U(c_t, h_t; C_\tau, H_\tau)$$

with properties analogous to those in (8), augmented by the conditions $U_H < 0$, $U_{HH} > 0$, and $U_{hH} > 0$. These conditions characterize H_τ as a relative reference point, where $U_{hH} > 0$ captures the idea that the marginal utility of labor (or leisure) increases with the aggregate benchmark.

Definition 2 (IPH(k_c, k_h) Preferences). *The preferences $U(c, h; C, H)$ are independent pairwise homogeneous of degrees k_c and k_h in the pairs (c, C) and (h, H) respectively if, for any scaling factors $\lambda_c, \lambda_h > 0$:*

$$U(\lambda_c c, \lambda_h h; \lambda_c C, \lambda_h H) = \lambda_c^{k_c} \lambda_h^{k_h} U(c, h; C, H) \quad (35)$$

Setting $\lambda_c = C_\tau^{-1}$ and $\lambda_h = H_\tau^{-1}$, we obtain the intensive formulation:

$$U(c_t, h_t; C_\tau, H_\tau) = C_\tau^{k_c} H_\tau^{k_h} u\left(\frac{c_t}{C_\tau}, \frac{h_t}{H_\tau}\right) \quad (36)$$

where $u\left(\frac{c_t}{C_\tau}, \frac{h_t}{H_\tau}\right) \equiv U\left(\frac{c_t}{C_\tau}, \frac{h_t}{H_\tau}; 1, 1\right)$ satisfies the regularity conditions:

$$u \in C^2, \quad u : \mathbb{R}_{++} \times [0, 1] \rightarrow \mathbb{R}, \quad u_c > 0, \quad u_{cc} < 0, \quad u_h < 0, \quad u_{hh} > 0 \quad (37)$$

Lemma 3. *The preferences $U(c, h; C, H)$ are $IPH(k_c, k_h)$ if and only if the marginal utility U_c is $IPH(k_c - 1, k_h)$ and the marginal utility U_h is $IPH(k_c, k_h - 1)$.*

Proof. We prove sufficiency and necessity separately:

- (a) The sufficiency (“if”) direction follows from Euler’s Theorem for homogeneous functions applied to each scaling pair. If the partial derivatives with respect to (c, C) are $IPH(k_c - 1, k_h)$ and those with respect to (h, H) are $IPH(k_c, k_h - 1)$, the total differential of U requires the primitive function to satisfy the $IPH(k_c, k_h)$ property.⁵
- (b) The necessity (“only if”) direction follows from differentiating the scaling identity $U(\lambda_c c, \lambda_h h; \lambda_c C, \lambda_h H) = \lambda_c^{k_c} \lambda_h^{k_h} U(c, h; C, H)$ with respect to c and h . Using the intensive formulation (36), we obtain:

$$U_c(c_t, h_t; C_\tau, H_\tau) = C_\tau^{k_c-1} H_\tau^{k_h} u_{\frac{c}{C}}\left(\frac{c_t}{C_\tau}, \frac{h_t}{H_\tau}\right) \quad (38)$$

$$U_h(c_t, h_t; C_\tau, H_\tau) = C_\tau^{k_c} H_\tau^{k_h-1} u_{\frac{h}{H}}\left(\frac{c_t}{C_\tau}, \frac{h_t}{H_\tau}\right) \quad (39)$$

where the arguments of the derivatives $u_{\frac{c}{C}}$ and $u_{\frac{h}{H}}$ depend only on the ratios c/C and h/H , rendering them homogeneous of degree $(0, 0)$. Thus, U_c scales by degree $(k_c - 1, k_h)$ and U_h by $(k_c, k_h - 1)$.

□

Lemma 4. *The ratio of $U_c(c, h; C, H)$ evaluated at two arbitrary allocations (c_t, h_t) and (c_{t+1}, h_{t+1}) satisfying $(c_{t+1}, h_{t+1}) = ((1 + \gamma)c_t, (1 + \gamma_h)h_t)$ and sharing common benchmarks C_τ and H_τ is independent pairwise homogeneous of degree zero, $IPH(0, 0)$, in the pairs (c, C) and (h, H) —meaning it remains scale-invariant under simultaneous independent scaling by factors $\lambda_c > 0$ in*

⁵Specifically, if U_c and U_C are $IPH(k_c - 1, k_h)$, Euler’s Theorem implies $k_c U = c U_c + C U_C$. Scaling the right-hand side by λ_c yields a factor of $\lambda_c^{k_c}$. Similarly, if U_h and U_H are $IPH(k_c, k_h - 1)$, then $k_h U = h U_h + H U_H$, and scaling by λ_h yields a factor of $\lambda_h^{k_h}$. Together, these satisfy $U(\lambda_c c, \lambda_h h; \lambda_c C, \lambda_h H) = \lambda_c^{k_c} \lambda_h^{k_h} U(c, h; C, H)$.

(c, C) and $\lambda_h > 0$ in (h, H) :

$$RMU = \frac{U_c(\lambda_c c_t, \lambda_h h_t; \lambda_c C_\tau, \lambda_h H_\tau)}{U_c(\lambda_c c_{t+1}, \lambda_h h_{t+1}; \lambda_c C_\tau, \lambda_h H_\tau)} = \frac{U_c(\lambda_c c_t, \lambda_h h_t; \lambda_c C_\tau, \lambda_h H_\tau)}{U_c(\lambda_c(1+\gamma)c_t, \lambda_h(1+\gamma_h)h_t; \lambda_c C_\tau, \lambda_h H_\tau)} = \frac{U_c(c_t, h_t; C_\tau, H_\tau)}{U_c((1+\gamma)c_t, (1+\gamma_h)h_t; C_\tau, H_\tau)}$$

if and only if the marginal utility of consumption, $U_c(c, h; C, H)$, is independent pairwise homogeneous of degrees $k_c - 1$ and k_h in the pairs (c, C) and (h, H) respectively; that is, $U_c(\lambda_c c, \lambda_h h; \lambda_c C, \lambda_h H) = \lambda_c^{k_c-1} \lambda_h^{k_h} U_c(c, h; C, H)$ for some constants $k_c - 1, k_h \in \mathbb{R}$.

Proof. We prove sufficiency and necessity separately:

- (a) U_c is IPH($k_c - 1, k_h$) $\implies$ RMU is IPH(0, 0) (sufficiency): This (the “if” direction) follows immediately from substituting an independent pairwise homogeneous function of degrees $k_c - 1$ and k_h into the ratio, where the joint scaling factors $\lambda_c^{k_c-1} \lambda_h^{k_h}$ cancel out from both the numerator and denominator.
- (b) RMU is IPH(0, 0) $\implies$ U_c is IPH($k_c - 1, k_h$) (necessity): To prove the “only if” direction, we assume that the ratio is scale-invariant, i.e., IPH(0, 0) under the independent scaling factors $\lambda_c > 0$ and $\lambda_h > 0$. Equating the baseline ratio with its scaled counterpart yields:

$$\frac{U_c(c_t, h_t; C_\tau, H_\tau)}{U_c((1+\gamma)c_t, (1+\gamma_h)h_t; C_\tau, H_\tau)} = \frac{U_c(\lambda_c c_t, \lambda_h h_t; \lambda_c C_\tau, \lambda_h H_\tau)}{U_c(\lambda_c(1+\gamma)c_t, \lambda_h(1+\gamma_h)h_t; \lambda_c C_\tau, \lambda_h H_\tau)}$$

This identity shows that shifting the individual allocation arguments along their respective BGP trends leaves the proportional scaling response of the marginal utility unchanged. Consequently, this scaling multiplier is independent of individual allocations and must depend strictly on λ_c and λ_h . Defining this function as $\phi(\lambda_c, \lambda_h)$, we can write

$$U_c(\lambda_c c_t, \lambda_h h_t; \lambda_c C_\tau, \lambda_h H_\tau) = \phi(\lambda_c, \lambda_h) U_c(c_t, h_t; C_\tau, H_\tau)$$

Sequential scaling by separate arbitrary bivariate expansion vectors $(\lambda_{c1}, \lambda_{h1})$ and $(\lambda_{c2}, \lambda_{h2})$ yields the bivariate version of Cauchy’s multiplicative functional equation $\phi(\lambda_{c1}\lambda_{c2}, \lambda_{h1}\lambda_{h2}) = \phi(\lambda_{c1}, \lambda_{h1})\phi(\lambda_{c2}, \lambda_{h2})$. Given that $U \in C^2$, $\phi(\lambda_c, \lambda_h)$ is continuous. Its unique non-trivial continuous solution is a joint power function product $\phi(\lambda_c, \lambda_h) = \lambda_c^{k_c-1} \lambda_h^{k_h}$ for some constants $k_c - 1, k_h \in \mathbb{R}$ (see Sydsæter et al. (2008)). Substituting this back provides $U_c(\lambda_c c_t, \lambda_h h_t; \lambda_c C_\tau, \lambda_h H_\tau) = \lambda_c^{k_c-1} \lambda_h^{k_h} U_c(c_t, h_t; C_\tau, H_\tau)$, completing the proof that U_c must be independent pairwise homogeneous of degrees $k_c - 1$ and k_h in the pairs (c, C) and (h, H) .

□

4.5.2. Balanced Growth with $IPH(k_c, k_h)$ Preferences

Taking the scaling factors C_τ and H_τ as given, an agent of age i at time t who enjoys a utility $U(c_t^i, h_t^i; C_\tau, H_\tau)$ chooses individual consumption c_t^i and labor supply h_t^i by solving:

$$\max_{\{c_t^i, h_t^i\}_{i=0}^I} \sum_{i=0}^I \beta^i U(c_t^i, h_t^i; C_\tau) \quad (40)$$

$$\text{s.t.} \quad \sum_{i=0}^I \frac{c_{t+i}^i}{(1+r)^i} = \sum_{i=0}^I \frac{w_{t+i} h_{t+i}^i}{(1+r)^i} + a_0 \quad (41)$$

given a_0 and a technologically determined growth rate. The intratemporal first-order condition for an agent of age i in period t is $U_h(c_t^i, h_t^i; C_\tau, H_\tau) = w_t U_c(c_t^i, h_t^i; C_\tau, H_\tau)$. The intertemporal first-order condition for an agent transitioning from age i at time t to age $i+1$ at $t+1$ is $U_c(c_t^i, h_t^i; C_\tau, H_\tau) = \beta(1+r)U_c(c_{t+1}^{i+1}, h_{t+1}^{i+1}; C_\tau, H_\tau)$.

Proposition 2. *The OLG economy described by (40)–(41) with preferences $U(c_t, h_t; C_\tau, H_\tau)$ admits a BGP with trending hours (32)–(34) if and only if the preferences are $IPH(k_c, k_h)$ in (c, C) and (h, H) ; that is,*

$$U(c_t, h_t; C_\tau, H_\tau) = C_\tau^{k_c} H_\tau^{k_h} u\left(\frac{c_t}{C_\tau}, \frac{h_t}{H_\tau}\right) \quad (42)$$

for any felicity function $u(\cdot, \cdot)$ satisfying the regularity conditions in (37).

Proof. (a) $IPH(k_c, k_h) \Rightarrow \text{BGP}$. The proof proceeds in four steps, analogous to Proposition 1:

1. *Intertemporal condition:* Plugging the $IPH(k_c, k_h)$ preferences (42) into the objective function (40) to derive the intertemporal condition and equating the intertemporal MRS for an agent of age i in period t to that of an agent of the same age i one cohort forward, yields, after cancelling scaling factors:

$$\frac{u_{\frac{c}{C}}\left(\frac{c_t^i}{C_\tau}, \frac{h_t^i}{H_\tau}\right)}{\beta u_{\frac{c}{C}}\left(\frac{c_{t+1}^{i+1}}{C_\tau}, \frac{h_{t+1}^{i+1}}{H_\tau}\right)} = \frac{u_{\frac{c}{C}}\left(\frac{c_{t+1}^i}{C_{\tau+1}}, \frac{h_{t+1}^i}{H_{\tau+1}}\right)}{\beta u_{\frac{c}{C}}\left(\frac{c_{t+2}^{i+1}}{C_{\tau+1}}, \frac{h_{t+2}^{i+1}}{H_{\tau+1}}\right)} = 1+r \quad (43)$$

Intertemporal $MRS(c_t^i, c_{t+1}^{i+1})$ Intertemporal $MRS(c_{t+1}^i, c_{t+2}^{i+1})$

That is, our objective is to show that the intertemporal MRS throughout the life cycle is invariant across cohorts.

2. *Stationarity of the intertemporal condition:* Plugging BGP constraints (32)–(34) into (43) to substitute in c_t^i, h_t^i, C_τ and H_τ , and canceling the growth terms $(1+\gamma)$ and $(1+\gamma_h)$

within arguments in the marginal utilities of the $\tau + 1$ cohort:

$$\frac{u_{\bar{c}}\left(\frac{c_t^i}{C_\tau}, \frac{h_t^i}{H_\tau}\right)}{\underbrace{\beta u_{\bar{c}}\left(\frac{(1+g_{i,c})(1+\gamma)c_t^i}{C_\tau}, \frac{(1+g_{i,h})(1+\gamma_h)h_t^i}{H_\tau}\right)}_{\text{Intertemporal } MRS(c_t^i, c_{t+1}^{i+1})}} = \frac{u_{\bar{c}}\left(\frac{c_t^i}{C_\tau}, \frac{h_t^i}{H_\tau}\right)}{\underbrace{\beta u_{\bar{c}}\left(\frac{(1+g_{i,c})(1+\gamma)c_t^i}{C_\tau}, \frac{(1+g_{i,h})(1+\gamma_h)h_t^i}{H_\tau}\right)}_{\text{Intertemporal } MRS(c_{t+1}^i, c_{t+2}^{i+1})}} = \text{Constant},$$

which is satisfied for all cohorts τ .

3. *Intratemporal condition:* Dividing the intratemporal *MRS* for an agent of age i in period t by that of an agent of the same age i in the subsequent cohort and rearranging, we obtain, after canceling scaling factors:

$$\frac{\underbrace{u_{\bar{c}}\left(\frac{c_t^i}{C_\tau}, \frac{h_t^i}{H_\tau}\right) \frac{1}{C_\tau}}_{\text{Intratemporal } MRS(c_t^i, h_t^i)}}{\underbrace{u_h\left(\frac{c_t^i}{C_\tau}, \frac{h_t^i}{H_\tau}\right) \frac{1}{H_\tau}}_{\text{Intratemporal } MRS(c_t^i, h_t^i)}} = \frac{w_{t+1}}{w_t} \frac{\underbrace{u_{\bar{c}}\left(\frac{c_{t+1}^i}{C_{\tau+1}}, \frac{h_{t+1}^i}{H_{\tau+1}}\right) \frac{1}{C_{\tau+1}}}_{\text{Intratemporal } MRS(c_{t+1}^i, h_{t+1}^i)}}{\underbrace{u_h\left(\frac{c_{t+1}^i}{C_{\tau+1}}, \frac{h_{t+1}^i}{H_{\tau+1}}\right) \frac{1}{H_{\tau+1}}}_{\text{Intratemporal } MRS(c_{t+1}^i, h_{t+1}^i)}} \quad (44)$$

Since wages grow, the intratemporal *MRS* must exactly offset this wage growth throughout the life cycle to ensure invariance across cohorts.

4. *Stationarity of the intratemporal condition:* Plugging BGP constraints to substitute in c_t^i , h_t^i , C_τ and H_τ into (44), canceling growth rates $(1+\gamma)$ and $(1+\gamma_h)$ inside the marginal utilities of the intratemporal *MRS* of cohort $\tau + 1$, and using the BGP condition that $(1+\gamma) = (1+\gamma_w)(1+\gamma_h)$:

$$\frac{\underbrace{u_{\bar{c}}\left(\frac{c_t^i}{C_\tau}, \frac{h_t^i}{H_\tau}\right)}_{\text{Intratemporal } MRS(c_t^i, h_t^i)}}{\underbrace{u_h\left(\frac{c_t^i}{C_\tau}, \frac{h_t^i}{H_\tau}\right)}_{\text{Intratemporal } MRS(c_t^i, h_t^i)}} = \frac{\underbrace{u_{\bar{c}}\left(\frac{c_t^i}{C_\tau}, \frac{h_t^i}{H_\tau}\right)}_{\text{Intratemporal } MRS(c_t^i, h_t^i)}}{\underbrace{u_h\left(\frac{c_t^i}{C_\tau}, \frac{h_t^i}{H_\tau}\right)}_{\text{Intratemporal } MRS(c_t^i, h_t^i)}} = \text{Constant}$$

Thus, if U is any regular felicity function with $IPH(k_c, k_h)$ in (c, C) and (h, H) , then the intertemporal and intratemporal conditions are stationary across all cohorts.

(b) BGP $\Rightarrow IPH(k_c, k_h)$.

1. *Intertemporal MRS must be stationary:* After equating the intertemporal *MRS* of an individual of age i in period t and that of an individual of the same age i one cohort

forward $(\tau + 1)$, the intertemporal condition along the BGP, is:

$$\begin{aligned} \frac{U_c(c_t^i, h_t^i; C_\tau, H_\tau)}{\beta U_c((1+\gamma)(1+g_{i,c})c_t^i, (1+\gamma_h)(1+g_{i,h})h_t^i; C_\tau, H_\tau)} &= \frac{U_c((1+\gamma)c_t^i, (1+\gamma_h)h_t^i; (1+\gamma)C_\tau, (1+\gamma_h)H_\tau)}{\beta U_c((1+\gamma)^2(1+g_{i,c})c_t^i, (1+\gamma_h)^2(1+g_{i,h})h_t^i; (1+\gamma)C_\tau, (1+\gamma_h)H_\tau)} \\ &= \frac{U_c((1+\gamma)c_t^i, (1+\gamma_h)h_t^i; (1+\gamma)C_\tau, (1+\gamma_h)H_\tau)}{\beta U_c((1+\gamma)(1+\gamma)(1+g_{i,c})c_t^i, (1+\gamma_h)(1+\gamma_h)(1+g_{i,h})h_t^i; (1+\gamma)C_\tau, (1+\gamma_h)H_\tau)} \\ &= 1 + r \end{aligned} \quad (45)$$

This equality implies that along the BGP, the intertemporal MRS must be IPH(0, 0), that is, independently scale-invariant with respect to the growing pairs (c, C) and (h, H) —requiring the consumption growth terms $(1 + \gamma)$ and the labor growth terms $(1 + \gamma_h)$ to cancel out uniformly. Consequently, the scaling multiplier is independent of individual allocations and depends strictly on the scaling factors, allowing us to express the marginal utility as:

$$U_c((1 + \gamma)c_t^i, (1 + \gamma_h)h_t^i; (1 + \gamma)C_\tau, (1 + \gamma_h)H_\tau) = \phi(1 + \gamma, 1 + \gamma_h)U_c(c_t^i, h_t^i; C_\tau, H_\tau)$$

which guarantees the stationarity of the intertemporal MRS. Since the discount factor β is a constant, the stationarity of the intertemporal MRS is equivalent to the stationarity of the ratio of marginal utilities across successive generations. Thus, by Lemma 4 (part (b)), this configuration necessitates that U_c is independent pairwise homogeneous of degrees $k_c - 1$ and k_h in (c, C) and (h, H) respectively, ensuring that the scaling vectors mapping cohort τ to $\tau + 1$ factor out uniformly from the marginal utilities to establish that $\phi(1 + \gamma, 1 + \gamma_h) = (1 + \gamma)^{k_c - 1}(1 + \gamma_h)^{k_h}$.

Thus, the stationarity of the intertemporal MRS implies that U_c is required to be IPH($k_c - 1, k_h$) in (c, C) and (h, H) , respectively.

2. *Intratemporal MRS must offset the grow rate of wages:* The intratemporal condition for two consecutive cohorts along the BGP is:

$$\frac{U_c(c_t^i, h_t^i; C_\tau, H_\tau)}{U_h(c_t^i, h_t^i; C_\tau, H_\tau)} = \frac{(1 + \gamma_w)w_t}{w_t} \frac{U_c((1 + \gamma)c_t^i, (1 + \gamma_h)h_t^i; (1 + \gamma)C_\tau, (1 + \gamma_h)H_\tau)}{U_h((1 + \gamma)c_t^i, (1 + \gamma_h)h_t^i; (1 + \gamma)C_\tau, (1 + \gamma_h)H_\tau)}$$

Using the necessary properties of U_c from the previous step:

$$\frac{U_c(c_t^i, h_t^i; C_\tau, H_\tau)}{U_h(c_t^i, h_t^i; C_\tau, H_\tau)} = (1 + \gamma_w)(1 + \gamma)^{k_c - 1}(1 + \gamma_h)^{k_h} \frac{U_c(c_t^i, h_t^i; C_\tau, H_\tau)}{U_h((1 + \gamma)c_t^i, (1 + \gamma_h)h_t^i; (1 + \gamma)C_\tau, (1 + \gamma_h)H_\tau)}$$

Then, in order to satisfy the BGP condition $(1 + \gamma) = (1 + \gamma_c) = (1 + \gamma_w)(1 + \gamma_h)$, we need U_h to be PH(k_c) in (c, C) and PH($k_h - 1$) in (h, H) . This implies that U_h must be

$IPH(k_c, k_h - 1)$, that is:

$$\begin{aligned} \frac{U_c(c_t^i, h^i; C_\tau, H_\tau)}{U_h(c_t^i, h^i; C_\tau, H_\tau)} &= (1 + \gamma_w) \frac{(1 + \gamma)^{k_c-1}}{(1 + \gamma)^{k_c}} \frac{(1 + \gamma_h)^{k_h}}{(1 + \gamma_h)^{k_h-1}} \frac{U_c(c_t^i, h^i; C_\tau, H_\tau)}{U_h(c_t^i, h^i; C_\tau, H_\tau)} \\ &= \frac{(1 + \gamma_w)(1 + \gamma_h)}{(1 + \gamma)} \frac{U_c(c_t^i, h^i; C_\tau, H_\tau)}{U_h(c_t^i, h^i; C_\tau, H_\tau)} = \frac{U_c(c_t^i, h^i; C_\tau, H_\tau)}{U_h(c_t^i, h^i; C_\tau, H_\tau)} \end{aligned}$$

Finally, collecting the joint requirement that U_c needs to be $IPH(k_c - 1, k_h)$ and U_h needs to be $IPH(k_c, k_h - 1)$ implies that the utility function U needs to be $IPH(k_c, k_h)$, as established in Lemma 3. □

Proposition 2 establishes that *independent pairwise homogeneity (IPH)* is the essential requirement for balanced growth when hours are trending. By allowing labor and consumption to scale with their respective aggregate benchmarks, this homogeneity property ensures BGP consistency for any regular felicity function, regardless of its specific functional form. The shape of the utility function is not an inherent property of balanced growth, regardless of whether hours trend or not.

4.5.3. Extension to Infinite-Horizon

In Appendix A.2, we extend the analysis to the infinite-horizon counterpart for the case of trending hours. This framework represents the limiting case of the OLG economy; although the first-order conditions for balanced growth are consistent across both models, the infinite-horizon environment requires a transversality condition to ensure a well-defined equilibrium path. We show that $IPH(k_c, k_h)$ is the essential preference property required to sustain balanced growth with trending labor supply.

4.5.4. The Current Paradigm: A Sub-Manifold of $IPH(k_c, k_h)$

We argue that the functional-form restrictions characterizing the Boppart and Krusell (2020) paradigm are not inherent requirements of balanced growth, but instead arise from the independent modeling choice of benchmark-independence. Crucially, the stationarity of the MRS between consumption and labor requires that the growth in the scaling factors C and H offsets the growth in the real wage. In the current paradigm, where these benchmarks are absent, this requirement imposes a strict link between the consumption curvature (k_c) and the labor curvature (k_h).

To see this, consider a BGP where hours shift at a rate relative to consumption growth such that the aggregate benchmarks are linked by $H \propto C^{-\frac{\gamma}{1-\gamma}}$. Substituting this relationship into the marginal utility of our $IPH(k_c, k_h)$ class yields:

$$U_c \propto C^{k_c-1-k_h \frac{\gamma}{1-\gamma}} u_{z_c}(z_c, z_h) \quad (46)$$

For benchmark-independence to hold, the marginal utility must be independent of the aggregate benchmark C . Taking the total derivative with respect to C and noting that $z_h = h/H$ depends on C via the BGP link ($z_h = hC^{\frac{\nu}{1-\nu}}$), the requirement $\frac{dU_c}{dC} = 0$ implies:

$$\frac{z_c u_{z_c z_c}}{u_{z_c}} = (k_c - 1) + \frac{\nu}{1 - \nu} \left(\frac{z_h u_{z_c z_h}}{u_{z_c}} - k_h \right) \quad (47)$$

Equation (47) demonstrates that when hours shift ($\nu \neq 0$), the curvature of the individual felicity function is strictly constrained by a linear combination of the cross-elasticity and the structural growth parameters.

To identify the unique functional form required, we solve the partial differential equation implied by $\frac{dU_c}{dC} = 0$ using the method of characteristics (see details in Appendix C). The characteristic equations are:

$$\frac{dz_c}{z_c} = \frac{dz_h}{-\frac{\nu}{1-\nu} z_h} = \frac{df}{[(k_c - 1) - k_h \frac{\nu}{1-\nu}]f} \quad (48)$$

where $f \equiv u_{z_c}$. Integrating the first two terms yields the characteristic coordinate $\xi = z_h z_c^{\frac{\nu}{1-\nu}}$. Solving for f as a function of ξ and substituting the BGP link $H \propto C^{-\frac{\nu}{1-\nu}}$ into the *IPH* formulation confirms that the aggregate benchmarks C and H mathematically cancel out, leaving the benchmark-independent marginal utility:

$$U_c(c, h) = c^{\bar{k}-1} \Phi(hc^{\frac{\nu}{1-\nu}}), \quad \bar{k} \equiv k_c - k_h \frac{\nu}{1 - \nu}, \quad (49)$$

where $\bar{k}$ is a composite of the two homogeneity degrees that reduces to k_c when hours are trendless ($\nu = 0$). Integrating with respect to c recovers the utility level $U(c, h) = c^{\bar{k}} \mathcal{F}(hc^{\frac{\nu}{1-\nu}})$. This confirms that the BK class is the unique sub-manifold where benchmark-independence is compatible with hours growth.

Crucially, (47) shows that benchmark-independence forces a rigid interlocking of the felicity function's curvatures: the consumption elasticity must be perfectly calibrated to the cross-elasticity and the long-run growth rate ν . Under this specific parameterization, the functional form effectively allows the aggregate trends C_t and H_t to mathematically vanish. This phenomenon creates the illusion that the BK preference shape is an inherent feature of balanced growth. In reality, this shape is merely the unique sub-manifold where preferences are forced to mimic the homogeneity of the environment to compensate for the absence of aggregate benchmarks. That is, the current paradigm's insistence on benchmark-independence breaks the natural “decoupling” between growth and functional forms in the general *IPH* class, forcing the curvature of preferences to fulfill a structural role naturally handled by the aggregate benchmarks in that more general environment. This dependency narrows the set of admissible preferences to the knife-edge case of Boppart and Krusell (2020), whereas *IPH* allows any regular preference structure to be consistent

with shifting-hours balanced growth.

4.6. The Decoupling of BGP and Functional Form

Our main theoretical result is that, in both OLG and infinite-horizon economies—and in the context of either trendless hours (King et al., 1988a) or trending hours (Boppart and Krusell, 2020)—consistency with balanced growth emerges from the property of *independent pairwise homogeneity* between individual choices and aggregate benchmarks, rather than from an inherent restriction on the functional form of preferences. We summarize this result in the following theorem:

Theorem 1. (Balanced Growth and IPH Preferences) *Consider an economy following a balanced growth path along which consumption (c) and labor income (wh) grow at rate $\gamma \geq 0$, labor supply (h) grows at rate γ_h , and interest rates (r) are constant. Any regular utility function $U(c, h; C_\tau, H_\tau)$ that is Independent Pairwise Homogeneous (IPH) of degrees k_c and k_h in the pairs (c, C_τ) and (h, H_τ) , respectively, admits a BGP for any reference benchmark $\tau \leq t$.*

Proof. See the proofs of Propositions 1–2 for OLG frameworks in Sections 4.2 and 4.5.2, respectively, and the proofs of Propositions 3–4 for their infinite-horizon counterparts in Appendix A. $\square$

Under *IPH*, the scaling of individual variables by their aggregate counterparts (C_τ, H_τ) ensures that the arguments of the intensive felicity function $u(c/C_\tau, h/H_\tau)$ remain constant along the BGP. Because the growth in these aggregate benchmarks performs the “lifting” required to balance the growing real wage and stationary interest rates, the intertemporal and intratemporal marginal rates of substitution remain stationary for any regular felicity function satisfying standard differentiability and concavity conditions. Further, while Theorem 1 establishes the generality of BGP consistency, it also explains why the existing literature—which additionally assumes benchmark-independent preferences—was forced to rely on strict functional forms. When benchmark-independence is imposed, the shape of the felicity function must be precisely chosen to cancel out the reference scales in $U(c, h; C, H)$. In this degenerate case, the “lifting” role of the aggregate benchmarks is lost, and the burden of growth consistency is transferred entirely onto the shape of the felicity function.

Corollary 1. (The benchmark-independent Sub-Manifold) *The class of benchmark-independent, BGP-consistent utility functions is a restricted sub-manifold of the $IPH(k_c, k_h)$ class. Specifically, a utility function $U(c, h; C, H) \in IPH(k_c, k_h)$ is benchmark-independent—that is, independent of the aggregate benchmarks (C, H) —if and only if it takes the form:*

$$U(c, h) = c^{\bar{k}} \mathcal{F}(hc^{\frac{\gamma}{1-\gamma}}) \quad (50)$$

where $\mathcal{F}(\cdot)$ is an arbitrary differentiable function, ν is the BGP parameter regulating the relative growth of aggregate hours and consumption, and $\bar{k} \equiv k_c - k_h \frac{\nu}{1-\nu}$ is a composite of the homogeneity degrees that coincides with k_c when hours are trendless ($\nu = 0$). This class of benchmark-independent preferences is the one defined in Boppart and Krusell (2020) for the BGP with trending hours ($\nu \neq 0$), and it subsumes the preferences of King *et al.* (1988a,b) as the special case with trendless hours ($\nu = 0$).

Proof. See Appendix C. □

We conclude that the necessity of specific functional forms is not an inherent property of balanced growth, but rather a mathematical consequence of assuming benchmark-independent felicity functions. The implied decoupling between the shape of preferences and BGP leaves open which member of the IPH class to use. Table 1 summarizes the *IPH* class for various selections of τ and k .⁶ Panel (a) illustrates the progression from the general case to our birth-level normalization ($\tau = t - i$), and ultimately to scale-invariant utility ($k = 0$). Panel (b) presents the corresponding preferences for infinite-horizon setups, while Panel (c) identifies the sub-manifold of standard benchmark-independent preferences.

5. The Degree of Pairwise Homogeneity and Related Preferences

Intuitively, the degree of *pairwise homogeneity* k between individual choices and their associated aggregate benchmark governs the relationship between long-run economic growth and the absolute level of lifetime utility, determining how much a cohort's welfare is boosted by being born into a wealthier era (a larger C_τ). If $k = 0$, the utility function is homogeneous of degree zero and welfare is purely relative; a cohort's utility depends only on its consumption relative to the benchmark, regardless of the absolute level of technology or wealth. If $k > 0$, future generations derive higher absolute utility from the same relative consumption bundles due to improved baseline conditions. Because C_τ is taken as given by the agent, the local curvature of utility is determined by the individual choice of c_t (or equivalently c_t/C_τ). However, as C_τ grows across cohorts, k effectively captures the trade-off between the benefits of aggregate growth and the diminishing marginal utility of that wealth. It is for this reason that, in the context of infinite-horizon models, the transversality condition (TVC) imposes limits on the admissible range of k . In Section 5.1, we show how the TVC defines this upper bound. In Section 5.2, we discuss the specific choice of $k = 0$, which makes the utility scale-invariant and satisfies the TVC for any growth rate. In Section 5.3, we show that alternative time-varying preferences in the literature consistent with a BGP are simply applications of the *pairwise homogeneity* principle under specific parameter values for k .

⁶In Appendix D, we further show that the class of IPH preferences can be re-expressed using the standards-of-living at the time of birth as a single benchmark.

Table 1: Balanced Growth and IPH Preferences: A Synthesis

<i>Independent Pairwise Homogeneous</i> Preferences:	Balanced Growth Path ($\gamma_c = \gamma_{wh} = \gamma$)	
	(a) $\gamma_h = 0$	(b) $\gamma_h \neq 0$
(a) OLG Models:		
▷ <i>General IPH Class:</i>		
$U(c, h; C, H) \in IPH(k_c, k_h)$ with $\tau \leq t$	$C_\tau^{k_c} u\left(\frac{c_t^i}{C_\tau}, h^i\right)$	$C_\tau^{k_c} H_\tau^{k_h} u\left(\frac{c_t^i}{C_\tau}, \frac{h_t^i}{H_\tau}\right)$
▷ <i>Birth-Level Normalization:</i>		
$U(c, h; C, H) \in IPH(k_c, k_h)$ with $\tau = t - i$	$C_{t-i}^{k_c} u\left(\frac{c_t^i}{C_{t-i}}, h^i\right)$	$C_{t-i}^{k_c} H_{t-i}^{k_h} u\left(\frac{c_t^i}{C_{t-i}}, \frac{h_t^i}{H_{t-i}}\right)$
▷ <i>Scale-Invariant Utility with Birth-Level Normalization:</i>		
$U(c, h; C, H) \in IPH(0, 0)$ with $\tau = t - i$	$u\left(\frac{c_t^i}{C_{t-i}}, h^i\right)$	$u\left(\frac{c_t^i}{C_{t-i}}, \frac{h_t^i}{H_{t-i}}\right)$
(b) Infinite-Horizon Models:		
▷ <i>General IPH Class:</i>		
$U(c, h; C, H) \in IPH(k_c, k_h)$ with $\tau \leq t$	$C_\tau^{k_c} u\left(\frac{c_t}{C_\tau}, h\right)$	$C_\tau^{k_c} H_\tau^{k_h} u\left(\frac{c_t}{C_\tau}, \frac{h_t}{H_\tau}\right)$
▷ <i>Scale-Invariant Utility:</i>		
$U(c, h; C, H) \in IPH(0, 0)$ with $\tau \leq t$	$u\left(\frac{c_t}{C_\tau}, h\right)$	$u\left(\frac{c_t}{C_\tau}, \frac{h_t}{H_\tau}\right)$
(c) A Sub-Manifold: benchmark-independent Preferences:		
▷ $U(c, h; C, H) = U(c, h)$	KPR Functional Form or BK Functional Form (with $\nu = 0$)	BK Functional Form (with $\nu \neq 0$)

Notes: The table classifies preferences based on the *Independent Pairwise Homogeneity* (IPH) property. In the OLG framework, i denotes the agent's age, and $\tau = t - i$ identifies the birth-year cohort. The *Scale-Invariant* case ($k_c = k_h = 0$) represents the purely relative status benchmark where utility levels are stationary along the BGP. Panel (c) identifies the sub-manifold where preferences are benchmark-independent, $U(c, h; C, H) = U(c, h)$. Within this sub-manifold, BGP consistency requires the functional forms of King et al. (1988a) (KPR) or Boppart and Krusell (2020) (BK). Specifically, BK preferences take the form $u(c, h) = \frac{1}{1-\sigma} [c \cdot \nu (hc^{\frac{\nu}{1-\nu}})]^{1-\sigma}$, where the parameter ν governs the relative strength of long-run income and substitution effects. As $\nu \rightarrow 0$, the income and substitution effects perfectly offset, and the BK form collapses to the KPR case, which implies trendless hours for any growth rate γ .

5.1. Transversality Condition

In the context of the OLG framework (Sections 4 and 4.5), the economy does not require an aggregate TVC for a BGP to exist. Although the benchmark C_τ grows across generations at rate γ , each cohort τ faces a finite planning horizon. Consequently, each agent's lifetime utility and

budget constraint remain bounded for any fixed C_τ , ensuring the individual optimization problem is well-defined regardless of long-run aggregate growth.

In contrast, within infinite-horizon models, the transversality condition (TVC) must hold for optimality and sufficiency. With trendless hours and $U(c_t, h; C_\tau)$, the TVC is $\lim_{t \rightarrow \infty} \beta^t U_c(c_t, h; C_\tau) a_t = 0$. Given $U_c \propto C_\tau^{k-1}$ and $a_t \propto (1 + \gamma)^t$, the TVC requires that along the BGP:

$$\beta(1 + \gamma)^k < 1 \quad (51)$$

That is, for any regular felicity function u , the TVC imposes an upper bound on the degree of pairwise homogeneity k to ensure that the lifetime budget converges. For example, in the current paradigm of benchmark-independent utility functions with trendless hours, the degree of homogeneity is linked to the curvature of the utility function by $k = 1 - \sigma$ (see Section 4.4). In this case, the TVC requires the well-known result $\beta(1 + \gamma)^{1-\sigma} < 1$ (King et al., 1988a,b). With our more general class of preferences that permits trends in hours by introducing the aggregate benchmark H_τ , the TVC is satisfied for the vector $k = (k_c, k_h)$ adhering to $\beta(1 + \gamma)^{k_c}(1 + \gamma_h)^{k_h} < 1$ (see Appendix A.2).

5.2. Scale-Invariant Utility: The Degree Zero

A simple and portable choice for the degree of pairwise homogeneity that trivially ensures the TVC (for discount factors $\beta < 1$) is $k = 0$. This choice offers several advantages. First, the specification provides significant simplicity; by setting $k = 0$, the model removes the need for preference parameters to “neutralize” the impact of growth on the interest rate, which simplifies to $(1 + r)^{-1} = \beta/(1 + \gamma)$. In other words, within our IPH framework, $k = 0$ ensures the scale invariance of utility levels. Under this parameterization, utility remains stationary as consumption grows because the negative externality of aggregate growth C_τ perfectly offsets the individual gains from private consumption c_t .

Second, scale-invariant utility aligns with ample research suggesting that baseline well-being does not grow indefinitely with absolute standards of living. The empirical validity of $k = 0$ rests largely on the so-called Easterlin Paradox (Easterlin, 1974), which observes that while richer individuals are happier at a point in time, average happiness remains stagnant over the long run as aggregate income rises (Clark et al., 2008). By the same token, the assumption of zero homogeneity precludes absolute growth effects on welfare, making it impossible to model scenarios where secular productivity gains make generations happier in an absolute sense. While Easterlin’s results were recently challenged by Stevenson and Wolfers (2008)—who found a positive link between GDP and subjective well-being across countries—Krueger (2008) notes that the time-series evidence remains ambiguous and often fails to show a significant long-term relationship between income growth and satisfaction. We interpret the collection of results as leaning towards scale-invariant utilities, although the empirical debate is not fully settled. Furthermore, Becker and

Rayo (2008) caution that happiness should be viewed as a “commodity” produced by the household rather than utility itself, suggesting that utility could still be rising due to improvements in non-happiness domains even if reported well-being appears flat.

5.3. Relation to Other Time-Varying Preference Structures

Here, we show that existing time-varying alternatives in the literature designed to deliver BGP-consistent preferences are, in fact, applications of the *pairwise homogeneity* principle with specific degrees of homogeneity. The spectrum of these approaches ranges from preference structures with consumption externalities (Section 5.3.1) and ad-hoc adjustments to labor disutility (Section 5.3.2), to more sophisticated internalizations, most notably the specification of Jaimovich and Rebelo (2009) (Section 5.3.3).

5.3.1. Preferences with Consumption Externalities

A central aspect of our analysis is the treatment of aggregate benchmarks as external. This positions our work within the literature on consumption externalities, which includes “keeping up with the Joneses” and “catching up with the Joneses” specifications. These are typically implemented via the timing of the aggregate benchmark, C_τ , where $\tau = t$ denotes a contemporaneous externality (keeping up) and $\tau = t - 1$ denotes a lagged one (catching up).⁷ In these models, as in our specification, the aggregate benchmark typically decreases an individual’s level of utility but increases the marginal utility of individual consumption. Effectively, the aggregate benchmark acts as a moving goalpost: as the economy grows, an agent must consume more simply to maintain a constant level of relative satisfaction. Our framework departs from these standard formulations by utilizing a birth-level normalization ($\tau = t - i$), which “freezes” the reference point at the moment of an agent’s entry into the economy.

We next examine whether the functional forms common in the consumption externalities literature satisfy the property of *pairwise homogeneity* required for balanced growth. We show that Joneses-style preferences, $U(c_t, C_\tau)$, typically assume specifications that are pairwise homogeneous in (c, C) . This makes them a subset of our general *IPH* class and, consequently, consistent with a BGP. Specifically, the most common specifications—multiplicative or additive in c and C —both satisfy *pairwise homogeneity*, albeit with different implications for the curvature of the utility function as the economy scales. Under the multiplicative structure (as in Galí (1994) and Abel (1990)):

$$U(c_t; C_\tau) = \frac{1}{1 - \sigma} \left(\frac{c_t}{C_\tau^\alpha} \right)^{1 - \sigma} \iff C_\tau^k u \left(\frac{c_t}{C_\tau} \right) = C_\tau^{(1 - \sigma)(1 - \alpha)} \frac{1}{1 - \sigma} \left(\frac{c_t}{C_\tau} \right)^{1 - \sigma} \quad (52)$$

⁷In this sense, agents do not internalize the impact of their consumption decisions on the aggregate benchmark. This distinguishes our framework from the internal habit specifications of Constantinides (1990) or Boldrin et al. (1997).

While Galí (1994) focuses on contemporaneous externalities ($\tau = t$), Abel (1990) allows for both internal and external effects. In the latter, the reference point is lagged aggregate consumption ($\tau = t - 1$) such that utility depends on $v_t = (c_{t-1}^D C_{t-1}^{1-D})^\theta$. When $D = 0$, the habit is purely external. As shown in (52), the multiplicative specification is *pairwise homogeneous* of degree $k = (1 - \sigma)(1 - \alpha)$ with the intensive felicity form $u(c/C) = \frac{1}{1-\sigma}(c/C)^{1-\sigma}$, thereby satisfying the requirements for a BGP.

Second, Campbell and Cochrane (1999) employs an additive externality:

$$U(c_t; C_t) = \frac{1}{1-\sigma}(c_t - \alpha C_t)^{1-\sigma} \iff C_\tau^k u\left(\frac{c_t}{C_\tau}\right) = C_\tau^{1-\sigma} \frac{1}{1-\sigma} \left(\frac{c_t}{C_\tau} - \alpha\right)^{1-\sigma}. \quad (53)$$

As shown, (53) is *pairwise homogeneous* of degree $k = 1 - \sigma$ in (c, C) . The intensive felicity formulation $u(c/C) = \frac{1}{1-\sigma}(c/C - \alpha)^{1-\sigma}$ thus satisfies the requirements for balanced growth.⁸

Thus, both the multiplicative (52) and additive (53) specifications, regardless of the choice of τ , fall within our general class of functions that satisfy *pairwise homogeneity* and are therefore consistent with balanced growth. Our main theoretical result emphasizes that the fundamental requirement for balanced growth is not the CRRA property, but rather *pairwise homogeneity*. Specifically, the relative risk aversion in the additive specification (53) of Campbell and Cochrane (1999) is not constant:

$$RRA = \frac{\sigma c_t}{c_t - \alpha C_\tau}.$$

That is, despite lacking constant curvature, this specification satisfies *pairwise homogeneity* and is consistent with a balanced growth path. This example highlights that pairwise homogeneity, rather than a CRRA functional form, is the essential condition for BGP consistency.

5.3.2. Ad-Hoc Adjustments and Scaling Factors

To circumvent the restrictive belief that only a narrow set of functional forms is consistent with a BGP, some authors incorporate exogenous TFP growth directly into the utility function. For example, Mertens and Ravn (2011) employ MaCurdy-type preferences with time-varying weights on labor disutility that grow at the rate of technological change, $Z_t = (1 + \gamma)^t$. While these adjustments could be regarded as ad-hoc mathematical devices to force BGP consistency, we show that they succeed precisely because they implicitly satisfy the requirements of the *IPH* class. Specifically, these preferences take the form:

$$U(c_t, h_t; Z_t) = \frac{1}{1-\sigma} c_t^{1-\sigma} - Z_t^{1-\sigma} \kappa \frac{h_t^{1+\frac{1}{\nu}}}{1+\frac{1}{\nu}} \iff U(c_t, h_t; Z_t) = Z_t^{1-\sigma} \left[\frac{1}{1-\sigma} \left(\frac{c_t}{Z_t}\right)^{1-\sigma} - \kappa \frac{h_t^{1+\frac{1}{\nu}}}{1+\frac{1}{\nu}} \right]. \quad (54)$$

⁸This additive form is also used by De la Croix (1996) in an OLG setting with birth-level normalization $\tau = t - i$.

As shown, these preferences are *pairwise homogeneous* in (c, Z) with a degree of homogeneity $k = 1 - \sigma$, thereby satisfying the requirements for a BGP with trendless hours. By explicitly incorporating the C_τ benchmark—grounded in psychological and sociological evidence—our approach provides a behavioral rationale for these practices. This transforms what were previously viewed as mathematical “tricks” into a preference structure that reflects how human expectations and reference points evolve alongside aggregate productivity.

5.3.3. The Jaimovich and Rebelo (2009) Preferences

Jaimovich and Rebelo (2009) (JR) seek preferences that are consistent with a BGP featuring trendless hours (where long-run income and substitution effects offset) while allowing substitution effects to dominate in the business cycle. They address the tension between short-run behavior and long-run consistency by introducing an endogenous scaling variable, X_t . The utility function and the law of motion for X_t are:

$$U(c_t, h_t; X_t) = \frac{1}{1 - \sigma} \left(c_t - \psi h_t^\theta X_t \right)^{1 - \sigma}, \quad X_t = c_t^\gamma X_{t-1}^{1 - \gamma}.$$

The parameter $\gamma \in [0, 1]$ governs the strength of the wealth effect on labor supply. As $\gamma \rightarrow 1$, X_t tracks c_t contemporaneously, and the utility collapses to the King et al. (1988a) (KPR) form. Conversely, as $\gamma \rightarrow 0$, X_t becomes constant, nesting the Greenwood et al. (1988) (GHH) form. Crucially, JR preferences are consistent with balanced growth if and only if $\gamma > 0$. Along a BGP where consumption grows at a constant rate, the recursive structure of X_t ensures the scaling factor eventually inherits that same growth rate. This allows the model to mimic GHH behavior in the short run (if γ is small) while satisfying BGP requirements in the long run.

Here, we highlight that the JR specification maps directly into our framework by satisfying *pairwise homogeneity* in (c, X) . By defining the intensive felicity function $u(c/X, h) = \frac{1}{1 - \sigma} (c/X - \psi h^\theta)^{1 - \sigma}$, we can rewrite the utility as:

$$U(c_t, h_t; X_t) = X_t^{1 - \sigma} \left[\frac{1}{1 - \sigma} \left(\frac{c_t}{X_t} - \psi h_t^\theta \right)^{1 - \sigma} \right].$$

This demonstrates that JR preferences are *pairwise homogeneous* of degree $k = 1 - \sigma$ in consumption and the scaling factor. In our framework, the structural “lifting” performed by the internal, recursive variable X_t is instead handled by the external benchmark C_τ . This achieves the same consistency as Jaimovich and Rebelo (2009)—allowing for flexible wealth effects—but avoids the analytical complexity of agents internalizing a law of motion for their own preference shifters. Notably, while the original JR formulation requires $\gamma > 0$ to maintain balanced growth, our use of an external benchmark ensures BGP consistency even in the limiting case where the short-run wealth effect is zero.

6. A Unified View: Identifying BGP-Consistent Preferences from the Micro Evidence

The primary implication of our theorem is a decoupling of preference parameters from long-run restrictions. In the standard paradigm with benchmark-independent preferences, BGP consistency acts as an “identifying restriction” that forces functional forms to match long-run patterns. This poses a tug-of-war: forcing functional forms to match long-run data implies that they may not simultaneously address other critical dimensions—e.g. microeconomic evidence, asset pricing and the business cycle. We organize this section around a single empirical anchor: micro estimates of the CRRA and the elasticity of labor supply as per Blundell et al. (2016). First, we show that our approach, birth-level normalization, simultaneously reconciles these micro estimates with the long-run behavior of hours, while freeing the curvature for asset pricing and breaking the Chetty (2006) bound. Second, we show that this decoupling is not merely permissive but *used* by the data. Specifically, our approach rationalizes the one first-order fact that a benchmark-independent felicity function cannot accommodate along a balanced growth path: the cross-cohort transformation of married women’s labor supply. The male Frisch margin, where benchmark-independent felicities and models with birth-level normalization coincide, serves as our control case; the household level, specifically women’s participation, is where the models diverge.

6.1. The Micro Anchor: A Stable Frisch and a Weakly Identified Curvature

The natural place to discipline the consumption curvature σ and the Frisch elasticity ψ is household panel data. Arguably, the most credible estimates are from Blundell et al. (2016), who estimate a collective life-cycle model of family labor supply. They obtain a male Frisch elasticity of $\psi \approx 0.59$ and leave the consumption curvature only weakly identified with a $\sigma \approx 2.5$. In our reproduction of their estimator, a cluster bootstrap places ψ in roughly $[0.32, 1.03]$ and σ in $[0.54, 8.3]$ at the 95% level. The question we are after is: are these micro estimates for ψ and σ elicited from household panel data consistent with long-run behavior? To frame our discussion, consider the following preferences:

$$U(c_{it}, h_{it}; C_\tau) = u\left(\frac{c_{it}}{C_\tau}, h_{it}\right) = \frac{1}{1-\sigma} \left(\frac{c_{it}}{C_\tau}\right)^{1-\sigma} - \frac{\kappa}{1+\frac{1}{\psi}} h_{it}^{1+\frac{1}{\psi}} \quad (55)$$

which are *pairwise homogeneous* in (c, C) with degree of homogeneity zero and assume birth-level normalization $\tau = t - i$.⁹ Setting the aggregate benchmark $C_\tau = 1$ eliminates the normal-

⁹That is,

$$U(c_{it}, h_{it}; C_\tau) = \frac{1}{1-\sigma} \left(\frac{c_{it}}{C_\tau}\right)^{1-\sigma} - \kappa \frac{h_{it}^{1+\frac{1}{\psi}}}{1+\frac{1}{\psi}} = C_\tau^0 \underbrace{\left(\frac{1}{1-\sigma} \left(\frac{c_{it}/C_\tau}{C_\tau/C_\tau}\right)^{1-\sigma} - \kappa \frac{h_{it}^{1+\frac{1}{\psi}}}{1+\frac{1}{\psi}} \right)}_{=u\left(\frac{c_{it}}{C_\tau}, h_{it}; \frac{C_\tau}{C_\tau}\right)} = C_\tau^0 \underbrace{\left(\frac{1}{1-\sigma} \left(\frac{c_{it}}{C_\tau}\right)^{1-\sigma} - \kappa \frac{h_{it}^{1+\frac{1}{\psi}}}{1+\frac{1}{\psi}} \right)}_{=u\left(\frac{c_{it}}{C_\tau}, h_{it}; 1\right)} = \frac{1}{1-\sigma} \left(\frac{c_{it}}{C_\tau}\right)^{1-\sigma} - \kappa \frac{h_{it}^{1+\frac{1}{\psi}}}{1+\frac{1}{\psi}} = u\left(\frac{c_{it}}{C_\tau}, h_{it}\right) \quad (56)$$

ization and returns to the standard paradigm with a benchmark-independent felicity function in MaCurdy (1981), as utilized in Blundell et al. (2016). Further, let the equilibrium time paths of $x_{it} = \{h_{it}, w_{it}, c_{it}\}$ be determined by an exogenous, technologically driven growth rate γ_x between periods t and $t + 1$, a life-cycle profile determined by growth rates $g_{t,x}$ between ages t and $t + 1$, and an idiosyncratic household shock ε_{jt} .¹⁰

In Blundell et al. (2016) the identification of ψ and σ relies on idiosyncratic unanticipated within-household variation over time. Using the intratemporal first-order condition for labor supply, $h_{it} = \kappa^{-\psi} w_{it}^{\psi} c_{it}^{-\sigma\psi} C_{\tau}^{\psi(1-\sigma)}$, and our defined time paths for x_{it} , the structural relation (in logs) across unanticipated shocks for a household j , $\ln x_{i+1,t+1} - \ln x_{i,t} - \mathbb{E}_j[\ln x_{i+1,t+1} - \ln x_{i,t}] = \Delta \varepsilon_{jt}^x$, can be written as:

$$\Delta \varepsilon_{jt}^h = \psi \Delta \varepsilon_{jt}^w - \sigma \psi \Delta \varepsilon_{jt}^c + u_{jt} \quad (57)$$

where u_{jt} is a normally distributed regression residual with mean zero and finite variance.

¹¹ Thus, a cohort's consumption-at-birth, C_{τ} , is differenced out. This implies that estimation (57)—or the Blundell et al. (2016) estimator for that matter—neither uses nor can reject birth-level normalization. Now, do these micro estimates for ψ and σ , though successfully elicited from household panel data, remain consistent with balanced growth?

To address this question, note that estimation (57) is based on a specific set of data moments: the within-household unexpected changes à la Blundell et al. (2016). But there are other moments generated by the model that can deliver identification. In particular, we are interested in moments that related long-run growth rates to the parameters of interest ψ and σ . This link emerges with the the across-cohort moments based on within-age variation (across time), $\mathbb{E}_j[\ln x_{i,t+1} - \ln x_{i,t}]$, which

with $u\left(\frac{c_t}{C_{\tau}}, h_t\right) = \frac{1}{1-\sigma} \left(\frac{c_t}{C_{\tau}}\right)^{1-\sigma} - \kappa \frac{h_t}{1+\frac{1}{\psi}}$. That is, under pairwise homogeneity of degree 0, $U(c_t, h_t; C_{\tau}) = C_{\tau}^0 u\left(\frac{c_t}{C_{\tau}}, h_t\right) = u\left(\frac{c_t}{C_{\tau}}, h_t\right)$.

¹⁰Specifically, for a household j with age i at time t we have $x_{jit} = x_{\tau,0}(1 + \gamma_x)^i \prod_{l=0}^{i-1} (1 + g_{l,x}) \exp(\varepsilon_{jt}) = x_{0,0}(1 + \gamma_x)^i \prod_{l=0}^{i-1} (1 + g_{l,x}) \exp(\varepsilon_{jt})$, where $(1 + \gamma_x)^i$ captures the BGP, $\prod_{l=0}^{i-1} (1 + g_{l,x})$ captures the cumulative lifecycle growth and $\exp(\varepsilon_{jt})$ is a multiplicative idiosyncratic shocks.

¹¹We present this estimation framework focusing on the FOC of labor, which follows MaCurdy (1981), solely for illustrative purposes; our objective is not to reassess the empirical estimates in Blundell et al. (2016). Naturally, a simple OLS specification would suffer from denominator bias and potential endogeneity if structural shocks are correlated with the error term. To circumvent these econometric issues, Blundell et al. (2016) implement GMM techniques, leveraging internal instruments—such as past lags of wages and hours—to instrument for current endogenous variables. Further, their estimation strategy relies on an econometric distinction between the nature of wage innovations: transitory wage shocks are utilized to identify the Frisch elasticity (ψ) because they isolate the pure intertemporal substitution effect without altering lifetime wealth, whereas permanent wage shocks are leveraged to pin down the consumption curvature (σ) and insurance parameters by shifting permanent income expectations. All these econometric issues lie beyond the scope of our analysis, which centers strictly on the consistency of the resulting micro estimates with balanced growth paths.

imply:

$$\gamma_h = \psi\gamma_w - \sigma\psi\gamma_c - \psi(1 - \sigma)(\ln C_{\tau+1} - \ln C_\tau) \quad (58)$$

where the far-right term captures the cohort changes emerging from the birth-level normalization in (55). Thus, since all these moments must hold at once, the procedure of only using within-household variation for identification (as in Blundell et al. (2016)) implicitly assumes that the preference parameters ψ and σ not only govern the within-cohort structural relationship of the idiosyncratic shocks in (57) but also the long-run growth in (58) independently. Is this assumption warranted? The answer lies in how we model preferences.

BGP-Inconsistency of benchmark-independent Preferences. Consider first benchmark-independent preferences—the standard paradigm—returning to the felicity function used in Blundell et al. (2016), where $C_\tau = 1$ in (55) for all cohorts, implying that cohort changes are null (that is, $\ln C_{\tau+1} - \ln C_\tau = 0$). In this case, if the preference parameters estimated from idiosyncratic shocks (57) yield $\sigma > 1$, then under steady wage and consumption growth ($\gamma_w = \gamma_c$), aggregate hours worked in (58) must decline along the balanced growth path (BGP). This directly violates the empirical long-run behavior of the U.S. economy, where aggregate hours per capita have remained remarkably stationary ($\gamma_h = 0$).¹² Thus, under the benchmark-independent paradigm, credible microeconomic estimates of ψ and σ are inconsistent with balanced growth.

The proposed solution (PH). In contrast, preferences that explicitly incorporate a birth-level normalization as (55) provide a reconciliation between the micro estimates and balanced growth. To see this, substituting $\ln C_{\tau+1} - \ln C_\tau = \gamma_c$ and the BGP restrictions $\gamma_w = \gamma_c$ into equation (58) implies that aggregate hours stationarity ($\gamma_h = 0$) holds for any combination of ψ and σ . This theoretical result follows from Proposition 1 in Section 4. Thus, our birth-level normalization ensures empirical micro estimates from Blundell et al. (2016) remain consistent with the stylized BGP facts that interest macroeconomists.


6.2. Breaking the Chetty (2006) Bounds on σ

Chetty (2006) identifies a fundamental conflict: in benchmark-independent models, the high risk aversion required to match asset-pricing data (Mehra and Prescott, 1985) implies a counterfactually massive long-run income effect, larger than the observed decline in hours observed across countries (except for the U.S Boppart and Krusell (2020)).

To illustrate this conflict in a dynamic setting, consider a 2-period model where an agent maximizes $\sum_{t=1,2} \beta^{t-1} u(c_t, h_t)$ with the felicity function (55) after setting $C_\tau = 1$ for all cohorts, i.e. the

¹²Boppart and Krusell (2020) document a generalized secular decline in hours across most developed countries over time, but show that the U.S. stands as a notable exception to this rule with trendless aggregate hours.

Figure 1: The benchmark-independent Felicity Constraint: Restricted Pairs (σ, ψ) for the Long-Run Decline of Hours



benchmark-independent preferences in (MaCurdy, 1981),

$$u(c_t, h_t) = \frac{c_t^{1-\sigma}}{1-\sigma} - \kappa \frac{h_t^{1+\frac{1}{\psi}}}{1+\frac{1}{\psi}} \quad (59)$$

subject to the intertemporal budget constraint $c_1 + \frac{c_2}{1+r} = w_1 h_1 + \frac{w_2 h_2}{1+r}$. In this setting, along a balanced growth path where c_t is proportional to $w_t h_t$, the labor supply at any date t is $\ln h_t = -\frac{1}{\frac{1}{\psi} + \sigma} \ln \kappa + \frac{1-\sigma}{\frac{1}{\psi} + \sigma} \ln w_t$. Hence, the uncompensated elasticity, which governs long-run effects, is:

$$\epsilon_{h,w}(\sigma, \psi) = \frac{\partial \ln h}{\partial \ln w} = \frac{1-\sigma}{\frac{1}{\psi} + \sigma}. \quad (60)$$

The Long-Run Decline in Hours. The preferences in (59) are inconsistent with a BGP with trendless hours unless $\sigma = 1$ (the KPR case). However, they can be consistent with a BGP where hours exhibit a trend: if $\sigma > 1$, income effects from wage growth dominate and hours decline.¹³ This question is quantitative. Following Boppart and Krusell (2020), an annual wage increase of 2% and an annual hours decline of 0.4% implies an uncompensated elasticity of $\epsilon_{h,w} = -0.2$. Figure 1 shows the contour set of (ψ, σ) pairs consistent with this uncompensated elasticity. For example, take the Blundell et al. (2016) point estimates at face value and ask whether they are consistent with the long-run behavior of hours under benchmark-independent felicity. Holding the Frisch elasticity at the estimated $\psi = 0.585$, benchmark-independent balanced growth at $\epsilon = -0.20$ requires $\sigma = 1.68$; the micro curvature $\sigma = 2.5$ instead implies $\epsilon = -0.356$, i.e. hours falling 0.71% per year—nearly twice the observed rate. At its point estimates the micro evidence is thus in

¹³Boppart and Krusell (2020) show that MaCurdy preferences are a specific case of BK preferences. The BK preferences nest the MaCurdy form when $\kappa = \frac{1}{1-\sigma}$ and $\nu = \frac{\sigma-1}{\frac{1}{\psi} + \sigma}$, so that ν equals the uncompensated elasticity (60).

Table 2: Bounded (σ, ψ) : Labor Collapse vs. Equity Premium

	(a) benchmark-independent Utility $U(c, h; C, H) = u(c, h)$			(b) IPH Preferences ($k = 0$) $U(c, h; C, H) = u(c/C, h/H)$	
	KPR	Calibration	Asset Prices	$\{\sigma, \psi\} = \text{Free}$	$\{\sigma, \psi\} = \{100, 0.5\}$
<i>Parameters:</i>					
Curvature (σ)	1.0	1.75	100.0	Free	100
Frisch Elasticity (ψ)	0.5	0.5	0.5	Free	0.5
<i>Long-Run Macro (100-Year):</i>					
Long-Run Elasticity of Labor	0.00	−0.20	−0.98	−0.20	−0.20
Annual Δh	0.00%	−0.40%	−1.96%	−0.40%	−0.40%
Remaining Hours (h_{100}/h_0)	100.0%	67.0%	13.8%	67.0%	67.0%
<i>Finance Dimension:</i>					
Equity Premium ($\approx \sigma \sigma_c^2$)	0.06%	0.11%	6.25%	Free	6.25%

Notes: Assumes 2% wage growth and 2.5% consumption volatility.

tension with the growth facts under benchmark-independent felicity: either the curvature is forced down to 1.68, contradicting the micro estimate, or keeping the micro estimate as elicited at $\sigma = 2.5$ hours should fall twice as fast as in the data. The tension is real at the point estimates but, given the weak identification of σ documented in Section 6.1, lies within sampling noise: the contour value $\sigma = 1.68$ sits inside the bootstrap interval $[0.54, 8.3]$. The point is not to reject any of these well-identified studies, but to show that birth normalization removes the need to take a stand.

The Asset Pricing Facts. To map the curvature parameter σ to asset-pricing facts, we introduce uncertainty in the second period following Mehra and Prescott (1985). Let second-period consumption growth be stochastic, such that $\ln(c_2/c_1) \sim N(g_c, \sigma_c^2)$. The Euler equation for a risky asset implies that the equity premium is approximately $\sigma \cdot \sigma_c^2$. This highlights the “tug-of-war”: a high σ is required to bridge the gap between low empirical consumption volatility ($\sigma_c \approx 0.025$) and the high historical equity premium ($\approx 6\%$), but this same σ dictates a counterfactual labor collapse.

Table 2 quantifies this tension. Panel (a) illustrates the failure of the benchmark-independent model to reconcile macro and finance facts. The KPR case ($\sigma = 1$) satisfies trendless hours but fails the finance dimension entirely. Conversely, targeting the historical equity premium requires a large $\sigma = 100$, which forces a labor supply collapse: hours decline by 1.96% annually, leaving only 13.8% of labor supply after a century.¹⁴ Clearly, the benchmark-independent felicity is

¹⁴The value $\sigma = 100$ is implausibly high as a literal coefficient of relative risk aversion; it is simply the curvature the i.i.d. calculation requires to match a 6% premium at 2.5% consumption volatility—the Mehra and Prescott (1985) equity-premium puzzle itself. We use it to make the labor cost stark, not to defend it.

squeezed between the premium and the labor trend: one cannot raise the former without steepening the latter.

The proposed solution (IPH). Under our framework with *independent pairwise homogeneity (IPH)* in (c, C) and (h, H) , this fundamental conflict disappears. To see this, let $U(c, h; C, H)$ be *IPH* and—without loss of generality—consider the scale-invariant case ($k_c = k_h = 0$) with a MaCurdy-style felicity function that extends that in (55) by incorporating an aggregate labor benchmark:¹⁵

$$U(c_t, h_t; C_\tau, H_\tau) = \frac{1}{1 - \sigma} \left(\frac{c_t}{C_\tau} \right)^{1 - \sigma} - \frac{\kappa}{1 + \frac{1}{\psi}} \left(\frac{h_t}{H_\tau} \right)^{1 + \frac{1}{\psi}}$$

where $U(c_t, h_t; C_\tau, H_\tau) = u(c_t/C_\tau, h_t/H_\tau)$ due to zero-degree homogeneity.¹⁶ The intratemporal first-order condition becomes:

$$\frac{u_h \left(\frac{c_t}{C_\tau}, \frac{h_t}{H_\tau} \right)}{u_c \left(\frac{c_t}{C_\tau}, \frac{h_t}{H_\tau} \right)} \equiv \frac{\kappa (h_t/H_\tau)^{1/\psi}}{(c_t/C_\tau)^{-\sigma}} = w_t \frac{H_\tau}{C_\tau}. \quad (61)$$

Along the BGP, both w_t and the benchmark ratio C_τ/H_τ grow at the same rate γ . Consequently, the RHS of (61) is invariant across cohorts. The aggregate income effect of growth is neutralized regardless of the value of σ . This theoretical result follows from Proposition 2 in Section 4.5. As shown in the final column of Table 2, we can set σ at whatever value asset pricing requires—100 to match the historical premium under the i.i.d. calculation, or the more moderate values implied by richer asset-pricing models, or the upper bound of 8.3 in the weakly identified curvature in Blundell et al. (2016)—together with $\psi = 0.5$ to match micro-elasticity estimates, all while maintaining the empirical labor trend. The *IPH* property thus decouples the micro estimates from long-run constraints, allowing them to be calibrated to other data moments such as the equity premium puzzle.

6.3. The Gender of Labor Supply: Birth Normalization in the Data

The reconciliation with the long-run behavior of hours and asset-pricing results in the previous subsection shows that the flexibility birth normalization grants is *useful*. This subsection demonstrates that it is also *used*: there is a first-order fact that a benchmark-independent felicity function cannot accommodate along a balanced growth path (BGP), yet which birth normalization successfully captures. This fact is the cross-cohort evolution of married women's labor supply. Our objective is not to redefine the BGP to accommodate long-run, gender-specific trends—the increase in married women's labor supply is explicitly non-linear (Doepke and Tertilt, 2016)—but

¹⁵The focus on $k_c = k_h = 0$ is innocuous, as the results hold for any degree of *IPH* (k_c, k_h); see Proposition 2.

¹⁶Note that setting $C_\tau = H_\tau = 1$ reverts the utility function to the benchmark-independent preferences in (59).

rather to test whether a birth-level normalization (econometrically, a cohort-varying preference anchor, H_{t-i} by gender) of preferences is demanded by the data to explain this cross-cohort rise.

A gendered loading on the work norm. The evolution of married women’s labor supply lives on the labor margin, making the relevant benchmark the gender-specific labor norm. In our framework, this norm is captured by the prevalence of gender-specific work in a cohort’s formative environment, denoted by H_{t-i}^j for $j \in \{m, f\}$, where m and f stand for male and female, respectively. We let member j ’s labor supply reference this norm with intensity θ_j :

$$u^j \left(\frac{c_t}{C_{t-i}}, \frac{h_{jt}}{(H_{j,t-i})^{\theta_j}} \right)$$

where setting $\theta_j = 0$ leaves member j ’s labor supply unnormalized.¹⁷

Empirical Strategy. A key empirical regularity over the last half-century is the stark asymmetry between genders: while female labor supply experienced a massive structural reallocation upward, male labor supply lacked any parallel generational shift. To capture this feature within our framework, we impose the structural identification restriction: $\theta_m = 0$. Imposing $\theta_m = 0$ is a parameter restriction stating that men are insulated from a same-gendered work norm; the contemporaneous average work effort of men that they observe when entering the labor market does not alter their benchmark utility on labor. labor market does not alter their baseline utility weight on labor. Consequently, their preference structure remains stable with respect to their own gendered benchmark, which, absent cross-spouse effects, collapses identically to the standard macro paradigm control case. This explains why the male Frisch elasticity estimated by Blundell et al. (2016) is remarkably stable over time. In contrast, married women’s labor supply—particularly the participation margin—is transformed across the very same cohorts, increasing over time. Accommodating such a structural increase implies $\theta_f > 0$, a key testable implication that we take to the data. Thus, our implementation focuses primarily on women. Our empirical strategy leverages a high-powered cross-cohort design made feasible by the scale of the Census. Here, female labor force participation is regressed on the work norm of the individual’s birth state and birth cohort, identifying the preference parameters via cross-state variation within specific cohorts.

Evidence. We use U.S. census microdata (Ruggles et al., 2025) on over one million married couples born between 1885 and 1955. We measure each spouse’s formative *work norm* H_{t-i} as the prevalence of female employment in their birth state during childhood, following the geographic,

¹⁷In Appendix F, we extend this framework to a specification that is independent pairwise homogeneous in (c, C) , and separately in (h_m, H_m) and (h_f, H_f) . The resulting gender-specific first order conditions are fully consistent with our empirical strategy, although we do not explicitly model the extensive margin of labor force participation. Further, we empirically focus on women throughout.

Table 3: Birth-Reference Loadings on Married Women’s Labor Force Participation

Specification	Coef. (×100)	(<i>t</i>)
<i>Panel A. Own formative work norm (1940–1980)</i>		
(1) Baseline	3.2	(6.6)
(2) + birth-state FE	1.7	(10.6)
(3) + cohort FE (cross-state within cohort)	1.55	(5.2)
(4) + marriage-market sex ratio	1.57	(5.3)
(5) + residence-state × year FE	0.95	(3.2)
<i>Panel B. Extended sample (1940–2019)</i>		
(6) Cohort FE	3.0	(8.2)
(7) + education-gap control	3.1	(8.4)
(8) Recent cohorts (born ≥ 1955)	0.88	(2.4)
(9) + residence-state × year FE	0.95	(2.1)
<i>Panel C. Husband’s formative work norm: cross-spouse (FFO)</i>		
(10) Baseline	3.2	(7.0)
(11) + birth-state and cohort FE	0.75	(2.5)
(12) Horse race: wife’s own exposure	0.8	(2.7)
(13) Horse race: husband’s exposure	1.8	(15.2)
<i>Panel D. Epidemiological: second-generation women</i>		
(14) Age controls only	1.6	(2.7)
(15) + US state and year FE	1.44	(4.3)
(16) + children	1.18	(3.0)

Notes: Each entry is the effect of a one-standard-deviation increase in the formative female-employment norm on the wife’s labor force participation (mean 0.36), in percentage points. Panels A–B use IPUMS USA married couples (Ruggles et al., 2025), 1940–1980 ($N \approx 1.23$ million couples; wife births 1885–1955); the formative work norm is the prevalence of female employment in the birth state during childhood. Panel C extends the sample to 1940–2019 ($N \approx 2.13$ million). Panel D assigns the labor force participation of first-generation (foreign-born) women by country of origin to US-born second-generation women through parental birthplace, 1940–1960 ($N \approx 77,500$, 47 origin countries). *t*-statistics (clustered by birth state in Panels A–C) in parentheses. See Appendix F.2 for further details on variable construction and empirical sample.

cohort-varying approach of Fogli and Veldkamp (2011). In our baseline specification without controls, a one-standard-deviation increase in the wife’s own formative norm is associated with a 3.2 percentage point increase in her adult labor force participation ($t = 6.6$); see row (1), Table 3. This loading remains highly significant at 1.7 percentage points when we introduce birth-state fixed effects ($t = 10.6$); see row (2), Table 3. Most cleanly, when we add cohort fixed effects to isolate the variations across birth states strictly within the same cohort, a wife’s participation rises with her own formative norm by 1.55 percentage points per standard deviation ($t = 5.2$), net of birth-state, cohort, and period effects; see row (3), Table 3.

A structural home and a magnitude. The reduced-form loading has a natural structural counterpart in the life-cycle female-participation model of Attanasio et al. (2008), where a married woman works when the gain from participation clears a threshold; birth-level normalization enters that threshold directly, so θ_f is a cohort shifter of the participation cutoff. Independently, the finding of Attanasio et al. (2018) that the aggregate labor-supply elasticity is not a structural constant but moves with the participation distribution empirically substantiates our theoretical result that the balanced growth does not restrict or discipline the Frisch elasticity.

To gauge the economic magnitude of these effects, we conduct a simple accounting exercise. Our preferred cross-state within-cohort specification (Table 3, Row 3) yields a significant and highly accurate coefficient of 1.55 percentage points. Evaluated at the baseline sample women participation mean of 0.36, this absolute shift translates into a relative extensive-margin response of $0.0155/0.36 \approx 0.043$. Accounting for intensive hours adjustments using an hours multiplier of 1.5 consistent with Attanasio et al. (2018) maps this empirical moment into a structural log-log elasticity baseline target of $0.043 \times 1.5 \approx 0.065$. Setting $j = f$ and $-j = m$ in the log-hours structural condition (F.1), the loading on the female work norm is $\theta_f(1 + \psi_f)$, where ψ_f is the Frisch elasticity for women. Borrowing the estimate of $\psi_f = 0.82$ from Blundell et al. (2016) and solving the structural identity $\theta_f \times 1.82 = 0.065$, we isolate a partial-equilibrium preference loading of $\theta_f \approx 0.036$. Finally, between cohorts born in 1915–25 and 1955–65, formative female-employment exposure rose about 1.5 standard deviations while married women’s participation rose 38 percentage points. Multiplying our localized loading by this historical shift in exposure implies that the birth-reference channel directly accounts for roughly 4–13% of that secular rise, or $\sim 6\%$ at our robust central estimate.¹⁸ This is a partial-equilibrium share; the dynamic social multiplier of Fogli and Veldkamp (2011), in which each generation learns from the participation of the one before, would amplify it.¹⁹

Summing up, we strongly reject the knife-edge case of $\theta_f = 0$, which defines the standard benchmark-independent paradigm wherein cohorts are entirely detached from their formative environments. Conversely, we also reject a literal full normalization of $\theta_f = 1$, which would counterfactually imply that adult labor supply is completely insulated from contemporaneous macro wage innovations. Instead, the interior baseline value of $\theta_f^{PE} \approx 0.036$ indicates that while women are highly responsive to current market signals, childhood exposure to local employment norms acts as a persistent cognitive anchor. Structurally, this modest value is sufficient to perform the

¹⁸Specifically, the central estimate multiplies the 1.5 standard deviation increase in formative exposure by the within-cohort cross-state coefficient of 1.55 percentage points to yield 2.33 percentage points, accounting for $2.33/38 = 6.1\%$ of the total historical rise. The lower bound uses the highly controlled estimate of 0.95 percentage points ($1.5 \times 0.95 = 1.43$ percentage points, or $1.43/38 = 3.8\%$), while the upper bound maps from the baseline specification of 3.2 percentage points ($1.5 \times 3.2 = 4.80$ percentage points, or $4.80/38 = 12.6\%$).

¹⁹Although not reported here, we have also conducted our analysis on PSID married couples (1968–2021). There, the rise in female labor supply appears—consistent with Attanasio et al. (2018)—on the extensive margin: a wife’s participation loads positively on her birth reference, while the intensive margin of already-working wives does not.

necessary “lifting” on the leisure margin, breaking the rigid parameter interlocking of the standard model and liberating the household preference specification to match micro panel data and macro growth trends simultaneously.

The reference binds on labor, not consumption. Birth-level normalization need not bind equally on every margin, and the evidence says it does not. A direct test of a *consumption* reference at birth—a Malmendier and Nagel (2011)-style design relating later behavior to cyclical consumption around birth—returns a robust null in our data: an effect on hours near +0.1% with a t -statistic close to zero once income-scarring controls are included, against a minimum detectable effect of about 2%. This is not a failure of the theory but a sharpening of it. Under $u(c/C_{t-i}, \cdot)$ relative risk aversion equals σ for every cohort regardless of C_{t-i} , so the consumption benchmark is nearly invisible at the household level; the formative reference for these cohorts instead operates through the *work* norm H_{t-i} —what counts as a normal amount of female work. The leisure-margin discipline balanced growth cannot supply is exactly the discipline the data do supply.

Ruling out the rival. The leading alternative hypothesis is that rising female bargaining power, rather than birth-referenced preferences, drives the secular shift in participation. In a collective household framework, the relative Pareto weight does not difference out; if it trends by cohort, it contaminates a naive estimate of θ_f . Indeed, our structural simulations show that a cohort-trending weight inflates the naive preference loading by 60% to 160% and, tellingly, drives the male placebo parameter (θ_m) away from zero. We address this potential contamination in three ways. First, the direct census loading survives the inclusion of the marriage-market sex ratio (leaving the coefficient essentially unchanged, $t \approx 5.3$) and residence-state-by-year fixed effects that absorb all state-level distribution factors simultaneously—including unilateral divorce laws, local female labor demand, and institutional policies. Under these strict controls, the coefficient falls slightly but remains positive and highly significant ($t \approx 3.2$; see Panel A in Table 3, Rows (4) and (5)); it is similarly robust to a predetermined spousal education-gap control and persists for more recent cohorts born after 1955 (see Panel B in Table 3). Second, the cross-spouse loading provides a preference-transmission signature that a Pareto weight shift alone cannot generate; the wife’s participation rises significantly with her husband’s formative exposure, replicating the cross-spouse “mothers and sons” pattern of Fernández et al. (2004), which no representative-agent or benchmark-independent model can rationalize (see Panel C in Table 3). Third, a clean $\theta_m \approx 0$ placebo effect is itself inconsistent with a dominant bargaining power channel. We therefore treat the direct θ_f estimate as our robust empirical core and the cross-spouse pattern as corroborating evidence. The estimates are economically modest but precisely identified, placing the discipline on preferences exactly where growth theory is silent: within the household utility structure and on the female participation margin.

It is preferences, not local economics. That the channel is operates through preferences rather than contemporaneous markets is shown most cleanly by an epidemiological design (Fernández et al., 2004). Among second-generation women (US-born with a foreign-born parent), participation rises with the female-work culture of their parents' country of origin—about +1.4 percentage points per standard deviation ($t \approx 4$)—holding US state and year fixed (see Panel D in Table 3). With US markets and institutions held constant, this isolates a transmitted preference component of the formative reference, independent of local labor markets.

The unified view. Summing up, the reconciliation of curvature with the hours decline, the freeing of σ for asset pricing, and the gendered participation result are then three readings of one device: the aggregate birth standard C_{t-i} , not the curvature of u , does the work balanced growth would otherwise force onto functional form. Where that work is invisible, that is, for the case of prime-age men, the benchmark coincides with the standard model ($\theta_m = 0$); where it is first-order, that is, for the case of women's participation across cohorts, the birth normalization is both necessary and, in the data, present.

7. Further Implications for Cross-Sectional and Business-Cycle Evidence

The reconciliation of microestimates of preference parameters (e.g., consumption curvature and the Frisch elasticity of labor supply) with the long-run decline of hours and asset pricing in a manner fully consistent with stylized balanced growth paths (Section 6.2) extends to other dimensions of the data. Here, we highlight two of these dimensions: cross-sectional evidence and business cycle behavior.

7.1. Reconciling Cross-Sectional Evidence with Balanced Growth

In the context of the cross-sectional behavior of hours and wages, the use of MaCurdy-type preferences (59), where consumption and leisure are separable, has been invoked to capture via $\sigma > 1$, or a dominant wealth effect, the low correlation between hours and wages observed in microdata (Pijoan-Mas, 2006; Heathcote et al., 2014), a phenomenon more recently discussed in Chang et al. (2020). This serves as a clear example where, to get cross-sectional evidence right, one is forced to give up BGP consistency under the standard paradigm where a high degree of curvature causes labor supply to trend toward zero as wages grow, making micro and macro evidence mutually exclusive, at least, in the context of the U.S., which has relatively trendless hours per capita.²⁰ In contrast, our framework resolves this tension by showing that micro-behavior and long-run growth are fundamentally decoupled. For example, in the case of trendless hours, setting the the degree of pairwise homogeneity in (c, C) to zero ($k = 0$), we have (55) and by delegating

²⁰Alternatively, Chang et al. (2020) show how high correlations in idiosyncratic preference and productivity shocks can generate the low correlation in wages and hours, even when using BGP consistent preferences that set $\sigma = 1$.

the lifting of secular trends to the aggregate benchmark C_τ , the felicity function is liberated from its structural role of neutralizing growth. Consequently, researchers can adopt a larger-than-one σ required to match the cross-sectional wage-hour correlations without jeopardizing the existence of a BGP.

7.2. Reconciling Business Cycle Evidence with Balanced Growth

A key aspect of business cycles is the relatively large fluctuations in aggregate employment or total hours. A successful approach to this question is that of indivisible labor which is consistent with a BGP with trendless hours (Hansen, 1985; Rogerson, 1988). In its simplest version, $u(c, h) = \ln c - \kappa h$, is a special case of KPR preferences. The intuition is that the aggregate elasticity of labor supply is large (∞ , in the limit), resulting in a linear disutility of labor. High aggregate elasticities are consistent with measures of aggregate hours that incorporate the extensive margin (Keane and Rogerson, 2012; Rogerson and Wallenius, 2009). However, these preferences are in conflict with BGP scenarios depicting a secular decline (Boppart and Krusell, 2020). Our approach makes indivisible labor preferences consistent with declining hours. For example, setting the degrees of pairwise homogeneity to $k_c = k_h = 0$, we have:

$$U(c, h; C, H) = u\left(\frac{c_t}{C_\tau}, \frac{h_t}{H_\tau}\right) = \ln\left(\frac{c_t}{C_\tau}\right) - \kappa \frac{h_t}{H_\tau}$$

which satisfies BGP requirements regardless of the long-run hours trend. Thus, our approach maintains the capacity of the indivisible labor mechanism to generate aggregate fluctuations while simultaneously allowing for trending hours.

Another prominent example is the preferences of Greenwood et al. (1988) (GHH), which remove income effects on labor supply. By eliminating these effects, GHH preferences are particularly suitable for environments where aggregate labor fluctuates significantly over the cycle. In the long run, however, the original GHH functional form implies a secular increase in hours that is at odds with a BGP characterized by either trendless or declining hours. In contrast, applying our approach with $k_c = k_h = 0$,

$$U(c_t, h_t; C_\tau, H_\tau) = u\left(\frac{c_t}{C_\tau}, \frac{h_t}{H_\tau}\right) = \frac{1}{1-\sigma} \left(\frac{c_t}{C_\tau} - \psi \left[\frac{h_t}{H_\tau} \right]^\theta \right)^{1-\sigma},$$

aligns GHH preferences with either trendless or declining hours while preserving the absence of income effects necessary to capture business cycle dynamics.

8. Conclusion

For decades, the prevailing view has been that models of economic growth require preferences with specific functional forms to be consistent with long-run trends. We show that this restriction

is not an inherent property of balanced growth, but rather a consequence of assuming benchmark-independent preferences. By introducing aggregate benchmarks grounded in psychology, sociology, and other social sciences, we show that *independent pairwise homogeneity* in preferences is the essential property for BGP consistency; under this property, any regular utility function can sustain a balanced growth path. Our approach decouples long-run growth from the shape of preferences. This decoupling liberates fundamental parameters, such as risk aversion and the Frisch elasticity of labor supply, from the rigid identifying restrictions of long-run trends, allowing researchers to target microeconomic, asset-pricing, business-cycle, or other evidence without compromising the stylized growth facts. In particular, we show that, anchored on a single set of micro estimates, the same device reconciles the consumption curvature with the secular decline in hours, frees that curvature for asset pricing, and—in census microdata on over a million married couples—rationalizes the gendered, early-formed transformation of women’s labor supply that a benchmark-independent felicity cannot carry along a balanced growth path. The freedom the theorem grants is the freedom the data demand.

References

- Abel, A. B. (1990). Asset prices under habit formation and catching up with the joneses. *The American Economic Review*, 80(2):38–42.
- Appadurai, A. (2004). The capacity to aspire: Culture and the terms of recognition. In Rao, V. and Walton, M., editors, *Culture and Public Action*. Stanford University Press.
- Attanasio, O., Levell, P., Low, H., and Sánchez-Marcos, V. (2018). Aggregating elasticities: Intensive and extensive margins of women’s labor supply. *Econometrica*, 86(6):2049–2082.
- Attanasio, O., Low, H., and Sánchez-Marcos, V. (2008). Explaining changes in female labor supply in a life-cycle model. *American Economic Review*, 98(4):1517–1542.
- Becker, G. S. (1996). *Accounting for Tastes*. Harvard University Press.
- Becker, G. S. and Rayo, L. (2008). Comment on Stevenson and Wolfers, “Economic growth and subjective well-being”. *Brookings Papers on Economic Activity*, 2008(1):88–95.
- Bick, A., Fuchs-Schündeln, N., and Lagakos, D. (2018). How do hours worked vary with income? cross-country evidence and implications. *American Economic Review*, 108(1):170–199.
- Bisin, A. and Verdier, T. (2001). The economics of cultural transmission and the dynamics of preferences. *Journal of Economic Theory*, 97(2):298–319.
- Blundell, R., Pistaferri, L., and Saporta-Eksten, I. (2016). Consumption inequality and family labor supply. *American Economic Review*, 106(2):387–435.
- Boldrin, M., Christiano, L. J., and Fisher, J. D. (1997). Habit persistence and asset returns in an exchange economy. *Macroeconomic Dynamics*, 1(2):312–332.
- Boppart, T. and Krusell, P. (2020). Labor supply in the past, present, and future: A balanced-growth perspective. *Journal of Political Economy*, 128(1):118–157.
- Campbell, J. Y. and Cochrane, J. H. (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):205–251.
- Chang, Y., Kim, S.-B., Kwon, K., and Rogerson, R. (2020). Aggregate labor supply in a heterogeneous-agent economy with idiosyncratic risk. *European Economic Review*, 124:103405.

- Chetty, R. (2006). A new method of estimating risk aversion. *American Economic Review*, 96(5):1821–1834.
- Clark, A. E., Frijters, P., and Gersttām, M. A. (2008). Relative income, happiness, and utility: An explanation for the easterlin paradox and other puzzles. *Journal of Economic Literature*, 46(1):95–144.
- Clausen, J. A. (1968). *Socialization and Society*. Little, Brown and Company.
- Constantinides, G. M. (1990). Habit formation: A resolution of the equity premium puzzle. *Journal of Political Economy*, 98(3):519–543.
- Cunha, F. and Heckman, J. J. (2007). The technology of skill formation. *American Economic Review*, 97(2):31–47.
- De la Croix, D. (1996). The dynamics of bequeathed tastes. *Economics Letters*, 51(1):89–96.
- Doepke, M. and Tertilt, M. (2016). Families in macroeconomics. In Taylor, J. B. and Uhlig, H., editors, *Handbook of Macroeconomics*, volume 2, pages 1789–1891. Elsevier.
- Duesenberry, J. S. (1949). *Income, Saving, and the Theory of Consumer Behavior*. Harvard University Press, Cambridge, MA.
- Easterlin, R. A. (1974). Does economic growth improve the human lot? Some empirical evidence. In David, P. A. and Reder, M. W., editors, *Nations and Households in Economic Growth: Essays in Honor of Moses Abramovitz*, pages 89–125. Academic Press, New York and London.
- Easterlin, R. A. (1980). *Birth and Fortune: The Impact of Numbers on Personal Welfare*. Basic Books.
- Easterlin, R. A., Pollak, R. A., and Wachter, M. L. (1980). Toward a more general model of fertility determination: Endogenous preferences and natural fertility. In *Population and Economic Change in Developing Countries*, pages 81–150. University of Chicago Press.
- Fernández, R., Fogli, A., and Olivetti, C. (2004). Mothers and sons: Preference formation and female labor force dynamics. *Quarterly Journal of Economics*, 119(4):1249–1299.
- Fogli, A. and Veldkamp, L. (2011). Nature or nurture? learning and the geography of female labor force participation. *Econometrica*, 79(4):1103–1138.
- Galí, J. (1994). Keeping up with the joneses: Consumption externalities, portfolio choice, and asset prices. *Journal of Money, Credit and Banking*, pages 1–8.
- Greenwood, J., Hercowitz, Z., and Huffman, G. W. (1988). Investment, capacity utilization, and the real business cycle. *The American Economic Review*, 78(3):402–417.
- Hansen, G. (1985). Indivisible labor and the business cycle. *Journal of Monetary Economics*, 16(3):309–327.
- Heathcote, J., Storesletten, K., and Violante, G. L. (2014). Consumption and labor supply with partial insurance: An analytical framework. *American Economic Review*, 104(7):2075–2126.
- Heckman, J. J. (2006). Skill formation and the economics of investing in disadvantaged children. *Science*, 312(5782):1900–1902.
- Jaimovich, N. and Rebelo, S. (2009). Can news about the future drive the business cycle? *American Economic Review*, 99(4):1097–1118.
- Keane, M. and Rogerson, R. (2012). Micro and macro labor supply elasticities: A distinctive heterogeneity approach. *Journal of Economic Literature*, 50(2):464–76.
- King, R. G., Plosser, C. I., and Rebelo, S. (1988a). Production, Growth and Business Cycles: I. The Basic Neoclassical Model. *Journal of Monetary Economics*, 21:195–232.
- King, R. G., Plosser, C. I., and Rebelo, S. (1988b). Production, Growth and Business Cycles: II. New Directions. *Journal of Monetary Economics*, 21:309–341.
- Krosnick, J. A. and Alwin, D. F. (1989). Aging and susceptibility to attitude change. *Journal of Personality and Social Psychology*, 57(3):416–425.
- Krueger, A. B. (2008). Comment on Stevenson and Wolfers, “Economic growth and subjective well-being”. *Brookings Papers on Economic Activity*, 2008(1):96–100.
- MaCurdy, T. E. (1981). An empirical model of labor supply in a life-cycle setting. *Journal of Political Economy*,

89(6):1059–1085.

- Malmendier, U. and Nagel, S. (2011). Depression babies: Do macroeconomic experiences affect risk taking? *The Quarterly Journal of Economics*, 126(1):373–416.
- McCrae, R. R. and Costa, P. T. (1994). The stability of personality: Observations and evaluations. *Current Directions in Psychological Science*, 3(6):173–175.
- Mehra, R. and Prescott, E. C. (1985). The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2):145–161.
- Mertens, K. and Ravn, M. O. (2011). Understanding the aggregate effects of anticipated and unanticipated tax policy shocks. *Review of Economic Dynamics*, 14(1):24–57.
- Pampel, F. C. and Peters, E. H. (1995). The social fabric and the effect of economic conditions on fertility. *Social Forces*, 73(4):1549–1571.
- Pijoan-Mas, J. (2006). Precautionary savings or working longer hours? *Review of Economic Dynamics*, 9(2):326–352.
- Pollak, R. A. (1976). Interdependent preferences. *The American Economic Review*, 66(3):309–320.
- Rogerson, R. (1988). Indivisible Labor, Lotteries and Equilibrium. *Journal of Monetary Economics*, 21(1):3–16.
- Rogerson, R. and Wallenius, J. (2009). Micro and macro elasticities in a life cycle model with taxes. *Journal of Economic Theory*, 144(6):2277–2292.
- Ruggles, S., Flood, S., Sobek, M., Backman, D., Cooper, G., Rivera Drew, J. A., Richards, S., Rodgers, R., Schroeder, J., and Williams, K. C. W. (2025). IPUMS USA: Version 16.0 [dataset].
- Stevenson, B. and Wolfers, J. (2008). Economic growth and subjective well-being: Reassessing the easterlin paradox. *Brookings Papers on Economic Activity*, 2008(1):1–87.
- Sydsæter, K., Hammond, P., Seierstad, A., and Strøm, A. (2008). *Further Mathematics for Economic Analysis*. Financial Times Prentice Hall, 2nd edition.
- Tarde, G. (1903). *The Laws of Imitation*. Henry Holt and Company, New York. First published in French as *Les lois de l'imitation* in 1890.
- Thompson, R. A., Hogan, A. E., and Clark, K. E. (2012). *The Development of Personality and Self-Regulation*. Oxford University Press.
- Wood, C., Littleton, K., and Sheehy, K. (2006). *Psychology and Education: A Developmental Analysis*. Routledge.

Online Appendix

Appendix A. Balanced Growth in Infinite-Horizon Models

In Appendix A.1 we show necessity and sufficiency of $PH(k)$ preferences for balanced growth in infinite horizon economies; the counterpart of the OLG environment in Section 4.2. In Appendix A.2 we extend the argument to infinite horizon economies with trending hours and $IPH(k_c, k_h)$ preferences, analogous to Section 4.5.2

Appendix A.1. Balanced Growth with $PH(k)$ Preferences in Infinite-Horizon Models

Consider an infinitely lived representative agent taking the initial aggregate benchmark C_τ with $\tau \leq t$ as given. The agent chooses sequences $\{c_t, h_t\}_{t=0}^\infty$ to solve:

$$\max_{\{c_t, h_t\}_{t=0}^\infty} \sum_{t=0}^\infty \beta^t U(c_t, h_t; C_\tau) \quad (\text{A.1})$$

$$\text{s.t.} \quad \sum_{t=0}^\infty \frac{c_t}{(1+r)^t} = \sum_{t=0}^\infty \frac{w_t h_t}{(1+r)^t} + a_0 \quad (\text{A.2})$$

The first-order conditions for this problem are:

$$U_c(c_t, h_t; C_\tau) = \beta(1+r)U_c(c_{t+1}, h_{t+1}; C_{\tau+1}) \quad (\text{A.3})$$

$$U_h(c_t, h_t; C_\tau) = w_t U_c(c_t, h_t; C_\tau) \quad (\text{A.4})$$

Proposition 3. *The infinite-horizon economy described by (A.1)–(A.4) with preferences $U(c_t, h_t; C_\tau)$ admits a BGP with trendless hours (5)–(7) if and only if the preferences are $PH(k)$ in (c, C) ; that is,*

$$U(c_t, h_t; C_\tau) = C_\tau^k u\left(\frac{c_t}{C_\tau}, h_t\right)$$

for any felicity function $u(\cdot, \cdot)$ satisfying the regularity conditions in (11).

Proof. (a) $PH(k) \Rightarrow \text{BGP}$.

1. *Intertemporal condition:* Under $PH(k)$, $U_c(c, h; C) = C^{k-1} u_c\left(\frac{c}{C}, h\right)$. Substituting into (A.3):

$$C_\tau^{k-1} u_c\left(\frac{c_t}{C_\tau}, h\right) = \beta(1+r) C_{\tau+1}^{k-1} u_c\left(\frac{c_{t+1}}{C_{\tau+1}}, h\right)$$

Along a BGP c and C grow at the same rate $(1+\gamma)$ and the intertemporal condition simplifies to:

$$C_\tau^{k-1} u_c\left(\frac{c_t}{C_\tau}, h\right) = \beta(1+r)(1+\gamma)^{k-1} C_\tau^{k-1} u_c\left(\frac{c_t}{C_\tau}, h\right)$$

which is satisfied by a constant interest rate $(1+r)^{-1} = \beta(1+\gamma)^{k-1}$.

2. *Intratemporal condition:* Under $PH(k)$ in (c, C) , $U_h = C_\tau^k u_h(c_t/C_\tau, h)$. Substituting into (A.4):

$$\underbrace{\frac{C_\tau^k u_h\left(\frac{c_t}{C_\tau}, h\right)}{C_\tau^{k-1} u_c\left(\frac{c_t}{C_\tau}, h\right)}}_{\text{Intratemporal } MRS(c_t, h_t)} = w_t \implies \underbrace{\frac{C_\tau u_h\left(\frac{c_t}{C_\tau}, h\right)}{u_c\left(\frac{c_t}{C_\tau}, h\right)}}_{\text{Intratemporal } MRS(c_t, h_t)} = w_t$$

where since both C_τ and w_t grow at the same rate γ , and u_h and u_c are stationary, then the growth of the intratemporal MRS exactly offsets that of wages.

(b) BGP $\Rightarrow$ PH(k).

1. *Intertemporal Necessity*: For the intertemporal MRS to be scale-invariant as c and C grow together, U_c must be homogeneous of some degree $k - 1$ in (c, C) . Thus, U_c is PH($k - 1$).
2. *Intratemporal Necessity*: From (A.4), $U_h/U_c = w_t$. Since w_t scales by $1 + \gamma$ and U_c scales by $(1 + \gamma)^{k-1}$, it follows that U_h must scale by $(1 + \gamma)^k$. By Lemma 1, U is PH(k).

□

Finally, in the context of infinite horizon models, for optimality and sufficiency, the transversality condition (TVC) $\lim_{t \rightarrow \infty} \beta^t U_c(c_t, h; C_\tau) a_t = 0$ must hold. Given $U_c \propto C_\tau^{k-1}$ and $a_t \propto (1 + \gamma)^t$, the TVC requires that along the BGP

$$\beta(1 + \gamma)^k < 1 \quad (\text{A.5})$$

That is, for any regular felicity function u , the TVC imposes an upper bound on the degree of pairwise homogeneity k to ensure that the lifetime budget converges; given a concave felicity function, this ensures the BGP is the unique global optimum. For example, in the context of the current paradigm with benchmark-independent utility functions with trendless hours that links the degree of homogeneity to the curvature of the utility function $k = 1 - \sigma$ (see Section 4.4), the TVC requires the well-known result that $\beta(1 + \gamma)^{1-\sigma} < 1$ (King et al., 1988a,b). With our more general class of preferences, the TVC is trivially satisfied for any k satisfying (A.5). In particular, setting $k = 0$, the TVC simply requires that $\beta < 1$.²¹

Appendix A.2. Balanced Growth with IPH(k_c, k_h) Preferences in Infinite-Horizon Models

Consider an infinitely lived representative agent taking the initial aggregate benchmarks C_τ and H_τ with $\tau \leq t$ as given. The agent chooses sequences $\{c_t, h_t\}_{t=0}^\infty$ to solve:

$$\max_{\{c_t, h_t\}_{t=0}^\infty} \sum_{t=0}^\infty \beta^t U(c_t, h_t; C_\tau, H_\tau) \quad (\text{A.6})$$

$$\text{s.t.} \quad \sum_{t=0}^\infty \frac{c_t}{(1+r)^t} = \sum_{t=0}^\infty \frac{w_t h_t}{(1+r)^t} + a_0 \quad (\text{A.7})$$

The first-order conditions for this problem are:

$$U_c(c_t, h_t; C_\tau, H_\tau) = \beta(1+r)U_c(c_{t+1}, h_{t+1}; C_{\tau+1}, H_{\tau+1}) \quad (\text{A.8})$$

$$U_h(c_t, h_t; C_\tau, H_\tau) = w_t U_c(c_t, h_t; C_\tau, H_\tau) \quad (\text{A.9})$$

Proposition 4. *The infinite-horizon economy described by (A.6)–(A.9) with preferences $U(c_t, h_t; C_\tau, H_\tau)$ admits a BGP with trending hours (32)–(34) if and only if the preferences are IPH(k_c, k_h) in (c, C) and (h, H) ; that is,*

$$U(c_t, h_t; C_\tau, H_\tau) = C_\tau^{k_c} H_\tau^{k_h} u\left(\frac{c_t}{C_\tau}, \frac{h_t}{H_\tau}\right)$$

for any felicity function $u(\cdot, \cdot)$ satisfying the regularity conditions.

Proof. (a) IPH(k_c, k_h) $\Rightarrow$ BGP.

²¹We did not discuss the transversality condition within the OLG framework in the main text because, unlike the infinite-horizon case, an OLG economy does not require an aggregate transversality condition for the existence of a BGP. Although the benchmark C_τ grows across generations at rate γ , each cohort τ faces a finite planning horizon. Consequently, each agent's lifetime utility and budget constraint remain bounded for any fixed C_τ , ensuring that the individual optimization problem is well-defined regardless of the long-run growth of the aggregate benchmark.

1. *Intertemporal condition:* Under *IPH*, $U_c(c, h; C, H) = C^{k_c-1} H^{k_h} u_{z_c}(z_c, z_h)$. Substituting into (A.8):

$$C_\tau^{k_c-1} H_\tau^{k_h} u_{z_c}(z_c, z_h) = \beta(1+r) C_{\tau+1}^{k_c-1} H_{\tau+1}^{k_h} u_{z_c}(z_c, z_h)$$

Along a BGP, C grows at rate $(1+\gamma)$ and H shifts at rate $(1+\gamma_h)$. The condition simplifies to:

$$1 = \beta(1+r)(1+\gamma)^{k_c-1} (1+\gamma_h)^{k_h}$$

which is satisfied by a constant interest rate $(1+r)^{-1} = \beta(1+\gamma)^{k_c-1} (1+\gamma_h)^{k_h}$.

2. *Intratemporal condition:* Under *IPH*, $U_h = C_\tau^{k_c} H_\tau^{k_h-1} u_{z_h}(z_c, z_h)$. Substituting into (A.9):

$$\frac{C_\tau^{k_c} H_\tau^{k_h-1} u_{z_h}(z_c, z_h)}{C_\tau^{k_c-1} H_\tau^{k_h} u_{z_c}(z_c, z_h)} = w_t \implies \frac{C_\tau u_{z_h}(z_c, z_h)}{H_\tau u_{z_c}(z_c, z_h)} = w_t$$

Since w_t and the benchmark ratio C_τ/H_τ both grow at the same rate $(1+\gamma)/(1+\gamma_h)$, and the arguments (z_c, z_h) are stationary, the growth of the intratemporal MRS exactly offsets that of wages.

(b) $\text{BGP} \implies \text{IPH}(k_c, k_h)$. For the intertemporal and intratemporal conditions to be scale-invariant as (c, C) and (h, H) grow along their respective trends, U_c and U_h must be independently homogeneous in their respective individual-aggregate pairs. This necessitates the *IPH*(k_c, k_h) structure. $\square$

Finally, in the context of infinite-horizon models, the transversality condition (TVC) $\lim_{t \rightarrow \infty} \beta^t U_c(c_t, h_t; C_\tau, H_\tau) a_t = 0$ must hold for optimality and sufficiency. Given that $U_c \propto C_\tau^{k_c-1} H_\tau^{k_h}$ and $a_t \propto (1+\gamma)^t$, the TVC requires that along the BGP:

$$\beta(1+\gamma)^{k_c} (1+\gamma_h)^{k_h} < 1 \quad (\text{A.10})$$

That is, for any regular felicity function u , the TVC imposes a joint upper bound on the degrees of homogeneity k_c and k_h to ensure that the lifetime budget converges. This condition is easily satisfied for a wide range of empirical parameters; for instance, setting $k_c = k_h = 0$ reduces the TVC to the standard requirement $\beta < 1$.

Appendix B. Implications of the benchmark-independence Constraint: Further Derivations

This appendix describes in more detail the functional form restrictions that emerge when benchmark-independence is imposed as a strict constraint on balanced growth, forcing $C_\tau = 1$ for all periods τ . In Proposition 1 we show that in our baseline framework with pairwise homogeneity (PH), the stationarity of the intertemporal marginal rate of substitution (MRS) across successive cohorts allows us to exploit an scaling function $\phi(\cdot)$, operating through the benchmark C_τ :

$$U_c((1+\gamma)c_t^i, h^i; (1+\gamma)C_\tau) = \phi(1+\gamma) U_c(c_t^i, h^i; C_\tau) \quad (\text{B.1})$$

Under this setup, Cauchy's multiplicative functional equation pins down the power-law curvature of the aggregate scaling factor, $\phi(1+\gamma) = (1+\gamma)^{k-1}$, leaving the local functional form of the individual felicity function entirely unrestricted.

Now, consider the benchmark-independent paradigm, which restricts the utility function such that $U(c, h; C_\tau) = U(c, h)$. Setting $C_\tau = 1$ for all periods effectively deactivates the aggregate benchmark. To preserve the stationarity of the intertemporal MRS along a BGP, the marginal utility of consumption must balance the growth factor natively:

$$\frac{U_c(c_t^i, h^i)}{\beta U_c((1+g_{i,c})(1+\gamma)c_t^i, h^i)} = \frac{U_c((1+\gamma)c_t^i, h^i)}{\beta U_c((1+g_{i,c})(1+\gamma)^2 c_t^i, h^i)} = 1+r \quad (\text{B.2})$$

Because the aggregate benchmarks cannot be updated to perform the structural “lifting,” this scale-invariance requirement can only be satisfied if shifting the absolute individual allocation argument maps directly to a proportional scaling response that is uniform across all initial levels:

$$U_c((1 + \gamma)c, h) = \Phi(1 + \gamma)U_c(c, h) \quad (\text{B.3})$$

Condition (B.3) reveals the structural trap forced by benchmark-independence: the primitive function U_c itself must swallow the aggregate growth factor natively. Crucially, because C_τ is fixed at unity, we cannot factor out an external scaling multiplier. Instead, Cauchy’s multiplicative functional equation must be applied directly to the primitive micro marginal utility function $U_c(\cdot, \cdot)$ itself with respect to its first argument, rather than to an external scaling factor, which ultimately requires the marginal utility to be a power function King et al. (1988a,b).

To see the exact mechanism that liberates preferences under our framework, compare where the scaling operator is applied across the two paradigms:

- Pairwise Homogeneity (our framework): Cauchy acts on an external, aggregate scaling filter $\phi(\cdot)$ that is independent of private choices, $U_c((1 + \gamma)c, h; (1 + \gamma)C_\tau) = \phi(1 + \gamma)U_c(c, h; C_\tau)$. Consequently, $\phi(\cdot)$ is forced to be a power function ($\phi(\lambda) = \lambda^{k-1}$), while the aggregate benchmark updates to absorb the macro scale. This leaves the internal felicity function $u(\cdot, \cdot)$ completely free to take any regular shape.
- benchmark-independence (the standard paradigm): Cauchy acts directly on the primitive micro marginal utility function $U_c(\cdot, \cdot)$, as shown in (B.3). Because there is no external benchmark to perform the lifting, the micro curvature of individual preferences is strictly dictated by macro growth trends.

Given that $U \in C^2$, the unique continuous non-trivial solution to this functional equation requires the primitive marginal utility to be an isoelastic power function in individual consumption:

$$U_c(c, h) = c^{k-1}v(h) \quad (\text{B.4})$$

Integrating this marginal utility with respect to c yields the standard constant relative risk aversion (CRRA) configuration:

$$U(c, h) = \begin{cases} \frac{c^{1-\sigma}}{1-\sigma}v(h) & \text{if } \sigma \neq 1 \\ \ln c + v(h) & \text{if } \sigma = 1 \end{cases} \quad (\text{B.5})$$

where $\sigma \equiv 1 - k$. This confirms that the functional-form discipline of the current paradigm is merely the unique sub-manifold where preferences are forced to mimic the homogeneity of the environment to compensate for the absence of aggregate benchmarks. Thus, the restriction to CRRA preferences is not an inherent property of balanced growth. Rather, it is an algebraic consequence of benchmark-independence. By setting $C_\tau = 1$, the standard paradigm forces the primitive preference curvature U_c to pull the structural weight that, under our more general pairwise homogeneous framework, is naturally handled by the cohort benchmark.

Appendix C. Recovering the Standard Paradigm (benchmark-independence) as Sub-Manifold of IPH: Detailed Derivations

In this appendix, we show the full derivation that formalizes the standard paradigm of benchmark-independent utility functions as a sub-manifold of the more general class of BGP-consistent *IPH* preferences. These derivations constitute the proof of Corollary 1. The results are discussed in the main text in Section 4.4 for the case of BGP with trendless hours (King et al., 1988a,b) and in Section 4.5.4 for the case of BGP with trending hours (Boppart and Krusell, 2020). The following proof proceeds to work out the more general case with trending hours, which subsumes the case of trendless hours.

Proof. Proof of Corollary 1. We proceed in five steps.

1. *The BGP link between scaling factors:* On a BGP, aggregate consumption C grows at the rate of technology γ . Following the same labeling for growth rate differences between consumption, wages and hours in Boppart and Krusell (2020), hours H shift at a rate relative to productivity growth governed by a parameter ν . Specifically, if the growth rate of consumption is γ , then the growth rate of hours is $\gamma_h = \gamma^{-\nu}$. In levels, this implies the structural relationship for the aggregate benchmarks:

$$H = AC^{-\frac{\nu}{1-\nu}} \quad (\text{C.1})$$

where A is a normalization constant.²²

2. *Total derivative of marginal utility:* Starting with the $IPH(k_c, k_h)$ marginal utility of consumption:

$$U_c(c, h; C, H) = C^{k_c-1} H^{k_h} u_{z_c}(z_c, z_h) \quad (\text{C.2})$$

where $z_c = c/C$ and $z_h = h/H$. Using the link in (C.1), we express U_c as a function of the single aggregate benchmark C :

$$U_c \propto C^{k_c-1-k_h\frac{\nu}{1-\nu}} u_{z_c}\left(\frac{c}{C}, hC^{\frac{\nu}{1-\nu}}\right) \quad (\text{C.3})$$

For benchmark-independence, we require $\frac{dU_c(c, h; C, H)}{dC} = \frac{dU_c(c, h)}{dC} \equiv 0$. Applying the product rule and the chain rule:

$$\begin{aligned} \frac{dU_c}{dC} &= \left(k_c - 1 - k_h \frac{\nu}{1-\nu}\right) C^{k_c-2-k_h\frac{\nu}{1-\nu}} u_{z_c} \\ &+ C^{k_c-1-k_h\frac{\nu}{1-\nu}} \left[u_{z_c z_c} \left(-\frac{c}{C^2}\right) + u_{z_c z_h} \left(\frac{\nu}{1-\nu} h C^{\frac{\nu}{1-\nu}-1}\right) \right] = 0 \end{aligned} \quad (\text{C.4})$$

Dividing by $C^{k_c-2-k_h\frac{\nu}{1-\nu}}$ and substituting $z_c = c/C$ and $z_h = hC^{\frac{\nu}{1-\nu}}$:

$$\left(k_c - 1 - k_h \frac{\nu}{1-\nu}\right) u_{z_c} - z_c u_{z_c z_c} + \frac{\nu}{1-\nu} z_h u_{z_c z_h} = 0 \quad (\text{C.5})$$

Rearranging and dividing by u_{z_c} yields the elasticity restriction in (47):

$$\frac{z_c u_{z_c z_c}}{u_{z_c}} = (k_c - 1) + \frac{\nu}{1-\nu} \left(\frac{z_h u_{z_c z_h}}{u_{z_c}} - k_h \right) \quad (\text{C.6})$$

3. *Solving the partial differential equation (PDE) via the method of characteristics:* Define $f(z_c, z_h) \equiv u_{z_c}$. The benchmark-independence requirement is a first-order linear PDE:

$$z_c f_{z_c} - \frac{\nu}{1-\nu} z_h f_{z_h} = \left[(k_c - 1) - k_h \frac{\nu}{1-\nu} \right] f \quad (\text{C.7})$$

The characteristic equations are:

$$\frac{dz_c}{z_c} = \frac{dz_h}{-\frac{\nu}{1-\nu} z_h} = \frac{df}{[(k_c - 1) - k_h \frac{\nu}{1-\nu}] f} \quad (\text{C.8})$$

First, solve for the characteristic path by integrating the first two terms:

$$\int \frac{1}{z_c} dz_c = -\frac{1-\nu}{\nu} \int \frac{1}{z_h} dz_h \implies \ln z_c = \ln(z_h^{-\frac{1-\nu}{\nu}}) + \ln C_1 \quad (\text{C.9})$$

²²This link ensures that the composite index $hC^{\frac{\nu}{1-\nu}}$ remains stationary on the BGP.

This gives the characteristic coordinate $\xi = z_h z_c^{\frac{\nu}{1-\nu}}$. Next, integrate the first and third terms:

$$\left[(k_c - 1) - k_h \frac{\nu}{1-\nu} \right] \int \frac{1}{z_c} dz_c = \int \frac{1}{f} df \implies f = C_2 z_c^{(k_c-1)-k_h \frac{\nu}{1-\nu}} \quad (\text{C.10})$$

Letting the constant C_2 be an arbitrary function $\Phi(\xi)$:

$$u_{z_c} = z_c^{(k_c-1)-k_h \frac{\nu}{1-\nu}} \Phi(z_h z_c^{\frac{\nu}{1-\nu}}) \quad (\text{C.11})$$

Substituting the BGP link $H \propto C^{-\frac{\nu}{1-\nu}}$ back into the *IPH* formulation, the aggregate benchmark cancels and the marginal utility is benchmark-independent. Writing it in levels, with $z_c = c/C$ and $z_h = hC^{\frac{\nu}{1-\nu}}$,

$$U_c(c, h) = c^{\bar{k}-1} \Phi\left(hc^{\frac{\nu}{1-\nu}}\right), \quad \bar{k} \equiv k_c - k_h \frac{\nu}{1-\nu}, \quad (\text{C.12})$$

so U_c depends on (c, h) only through the consumption level c and the stationary composite $hc^{\frac{\nu}{1-\nu}}$. The exponent inherited from the characteristic solution, $(k_c - 1) - k_h \frac{\nu}{1-\nu}$, is precisely $\bar{k} - 1$.²³

4. *Recovering the utility function:* We integrate $U_c(c, h) = c^{\bar{k}-1} \Phi(hc^{\frac{\nu}{1-\nu}})$ with respect to c at fixed composite $x \equiv hc^{\frac{\nu}{1-\nu}}$. Because the integrand is the product of the power $c^{\bar{k}-1}$ and a function of x alone, integration regroups the powers of c into the single exponent $\bar{k}$, yielding

$$U(c, h) = c^{\bar{k}} \mathcal{F}\left(hc^{\frac{\nu}{1-\nu}}\right), \quad (\text{C.13})$$

for an arbitrary differentiable $\mathcal{F}$; any additive function of h alone is incompatible with *IPH* homogeneity and is therefore zero.

5. *Mapping to the standard paradigm functional forms:* The general solution $U(c, h) = c^{\bar{k}} \mathcal{F}(hc^{\frac{\nu}{1-\nu}})$, with the composite exponent $\bar{k} = k_c - k_h \frac{\nu}{1-\nu}$, defines the broad class of preferences compatible with benchmark-independence, which are identical to those of Boppart and Krusell (2020). To see this, set $\bar{k} = 1 - \sigma$ and re-parameterize the arbitrary function as $\mathcal{F}(\xi) \equiv [v(\xi)]^{1-\sigma}/(1-\sigma)$. Substituting these into the integrated result, and distributing the exponent allows us to group c and $v(\cdot)$ within the same power transformation:

$$U(c, h) = \frac{\left[c \cdot v(hc^{\frac{\nu}{1-\nu}}) \right]^{1-\sigma}}{1-\sigma} \quad (\text{C.14})$$

Finally, note that by setting $\nu = 0$ (no trend in hours), the utility function in (C.14) reverts to the King et al. (1988a,b) preferences. Specifically, the homogeneity degree k_c dictates the functional form within this class: for $k_c \neq 0$, we recover the standard power utility $U(c, h) = \frac{c^{k_c}}{k_c} v(h)$, while the limit as $k_c \rightarrow 0$ recovers the log-separable case $U(c, h) = \ln c + v(h)$. This confirms that the *IPH* framework provides a unified structure that subsumes the entire evolution of the current paradigm.

□

Appendix D. Re-Expressing IPH Preferences with a Single Benchmark

The IPH class of preferences $u(c/C_\tau, h/H_\tau)$, which supports a BGP for any regular felicity u , can be re-expressed with a single preference benchmark—the cohort's standard of living C_τ —by letting the cohort's birth wage w_τ

²³On the BK sub-manifold the household's optimality conditions hard-wire the link between the curvatures and the hours trend: the stationarity of the labor-supply condition along the BGP, $(1 + \gamma_w)(1 + \gamma_h) = (1 + \gamma)$, requires (solving the FOCs without aggregate benchmarks) that the rate at which hours shift satisfy $k_h + \nu(k_c - k_h) = 0$, i.e. $\nu = -k_h/(k_c - k_h)$, so that the intertemporal and intratemporal marginal rates of substitution remain consistent with a constant interest rate and a growing wage.

(a market price, not a preference parameter) supply the second scaling. The felicity becomes $U(c, h; C_\tau, w_\tau) = C_\tau^k u(c/C_\tau, w_\tau h/C_\tau)$, which is exactly IPH under the identification $H_\tau \equiv C_\tau/w_\tau$. To see this, let C_τ denote per-capita consumption at the cohort's birth (its standard of living) and w_τ the wage prevailing then. Then, specify the felicity

$$U(c_t, h_t; C_\tau, w_\tau) = C_\tau^k u\left(\frac{c_t}{C_\tau}, \frac{w_\tau h_t}{C_\tau}\right), \quad (\text{D.1})$$

whose two arguments are consumption relative to the living standard, c_t/C_τ , and labor income at the birth wage relative to the living standard, $w_\tau h_t/C_\tau$ —the share of the cohort's standard of living that a member's work would command at the wages of its youth.

It is a single-benchmark, leisure-central member of the IPH class. Define $H_\tau \equiv C_\tau/w_\tau$. Then $w_\tau h_t/C_\tau = h_t/H_\tau$, and (D.1) is exactly the $IPH(k, 0)$ preference $C_\tau^k u(c_t/C_\tau, h_t/H_\tau)$, which by Theorem 1 admits a balanced growth path for any regular felicity u . The labor benchmark is therefore not an independent free object: along the BGP it is pinned by the living standard and the wage, $H_\tau = C_\tau/w_\tau$, and grows at the hours rate $(1 + \gamma_h) = (1 + \gamma)/(1 + \gamma_w)$. A single *preference* benchmark—the standard of living—thus does the work of both, with the market wage (a price agents already face, not an estimated parameter) supplying the second scaling. The leisure margin is built in: the second argument is the labor/earnings margin measured against the living standard, and it is here that the Frisch elasticity is freed.

Balanced growth for any u , with a transparent transversality bound. Substituting (D.1) into the household problem, the intratemporal condition reduces to

$$\frac{u_{z_h}}{u_{z_c}} = \frac{w_t}{w_\tau} = (1 + \gamma_w)^i \quad \text{for an age-}i \text{ agent,}$$

which depends on age but not on the cohort, so it is stationary across generations for any regular u ; the intertemporal condition is stationary for the same reason, since the common factor C_τ^{k-1} cancels in the within-cohort Euler and the arguments $(c/C_\tau, w_\tau h/C_\tau)$ are themselves stationary (this is the $IPH(k, 0)$ case of Proposition 2). In the infinite-horizon counterpart the transversality condition is simply

$$\beta(1 + \gamma)^k < 1, \quad (\text{D.2})$$

the *IPH* bound $\beta(1 + \gamma)^{k_c}(1 + \gamma_h)^{k_h} < 1$ at $(k_c, k_h) = (k, 0)$; the hours trend does not enter, because with u general the wage cannot be factored out as an outer power of H_τ . The scale-invariant choice $k = 0$ satisfies it for any growth rate.

It nests the paradigm in both regimes. With trendless hours ($\gamma_h = 0$, so $w_\tau \propto C_\tau$) the second argument reduces to hours alone and (D.1) collapses to the consumption-only $PH(k)$ form $C_\tau^k u(c_t/C_\tau, h_t)$, recovering King et al. (1988a) for the CRRA shape; with declining hours it reproduces the Boppart and Krusell (2020) income-effect labor supply—for any regular u , not only the knife-edge functional forms.

Appendix E. Reproducing the BPS Estimates and Their Sampling Uncertainty

The micro anchor of Section 6.1 is the family labor-supply GMM of Blundell et al. (2016). We reproduce their Table 6, column 2 specification on the same PSID sample, porting the estimator to a stand-alone implementation and cross-checking it against the original Stata code; the two agree to within 0.4% on every elasticity once the common- N weighting convention of the original is matched (the criterion normalizes every moment by the common complete-data sample size rather than each moment's own count). We recover a male Frisch elasticity $\psi = 0.585$ (analytic standard error 0.075), a consumption-price elasticity $\eta_{c,p} = 0.400$, and hence a curvature $\sigma = 1/\eta_{c,p} \approx 2.5$.

The curvature is weakly identified. Resampling households with a cluster bootstrap (220 replications) yields 95% intervals $\psi \in [0.32, 1.03]$ and $\sigma = 1/\eta_{c,p} \in [0.54, 8.3]$ (median $\sigma \approx 2.17$). The balanced-growth contour value implied by the observed hours decline at $\psi = 0.585$, namely $\sigma = 1.68$, lies well inside this interval. The point-estimate tension between the micro curvature and the secular hours decline discussed in Section 6.1 is therefore within sampling

noise—which is why birth-level normalization, by reconciling any (σ, ψ) with the trend, removes the need to take a stand on where in this wide interval the curvature lies.

Appendix F. Evidence of Gender Work Norms: A Framework for Hypothesis Testing, Data and Empirical Sample

Appendix F.1. The Analytical Framework

To capture the joint labor supply dynamics of dual-earner households consistent with balanced growth, we use preferences that embed both cohort-specific consumption standards and gender-specific formative work norms. The joint contemporaneous utility block for a household consisting of a female (f) and a male (m) earner is:

$$U\left(\frac{c_t}{C_{t-i}}, \mathbf{h}_t; \mathbf{H}_{t-i}\right) = \Lambda(\mathbf{H}_{t-i}) \left[\frac{1}{1-\sigma} \left(\frac{c_t}{C_{t-i}}\right)^{1-\sigma} - \sum_{j \in \{f, m\}} \Omega_{-j}(\mathbf{H}_{t-i}) \frac{\kappa_j}{1 + \frac{1}{\psi_j}} \left(\frac{h_{jt}}{H_{j,t-i}^{\theta_j}}\right)^{1 + \frac{1}{\psi_j}} \right]$$

where $\mathbf{h}_t \equiv \{h_{mt}, h_{ft}\}$ represents the vector of individual spousal labor supplies, and $\mathbf{H}_{t-i} \equiv \{H_{m,t-i}, H_{f,t-i}\}$ denotes the vector of gender-specific formative labor norms inherited by the cohort. The parameter ψ_j characterizes the gender-specific micro Frisch elasticity parameter.

The preference specification satisfies *independent pairwise homogeneity* (IPH) in (c, C) and (h_j, H_j) separately by gender with

$$\Lambda(\mathbf{H}_{t-i}) \equiv \prod_{j \in \{f, m\}} H_{j,t-i}^{k_j}$$

where the degree of pairwise homogeneity is gender specific and equal to $k_j = (1 - \theta_j) \left(1 + \frac{1}{\psi_j}\right)$. This IPH property allows for trends in gender-specific labor supply regardless of the curvature parameter σ and elasticities ψ_j . The cross-gender labor supply dependencies are captured via the additional component in the disutility weight $\Omega_{-j}(\mathbf{H}_{t-i}) \equiv H_{-j,t-i}^{k_{-j}}$, where the notation $-j$ indexes the counterpart gender in the dual-earner couple (i.e., $-j = f$ if $j = m$, and vice versa).

The implied intertemporal first order condition for h_{jt} is,

$$\begin{aligned} h_{jt}^{\frac{1}{\psi_j}} &= \frac{1}{\kappa_j} \left(\frac{c_t}{C_{t-i}}\right)^{-\sigma} \frac{w_{jt}}{C_{t-i}} H_{-j,t-i}^{-k_{-j}} H_{j,t-i}^{\theta_j \left(1 + \frac{1}{\psi_j}\right)} \\ &= \frac{1}{\kappa_j} w_{jt} c_t^{-\sigma} C_{t-i}^{\sigma-1} H_{-j,t-i}^{(\theta_{-j}-1) \left(1 + \frac{1}{\psi_{-j}}\right)} H_{j,t-i}^{\theta_j \left(1 + \frac{1}{\psi_j}\right)} \end{aligned}$$

Isolating h_{jt} and taking logs:

$$\ln h_{jt} = -\psi_j \ln \kappa_j + \psi_j \ln w_{jt} - \psi_j \sigma \ln c_t + \psi_j (\sigma - 1) \ln C_{t-i} + \underbrace{\theta_j (\psi_j + 1) \ln H_{j,t-i}}_{\text{Own Work Norm}} - \underbrace{(1 - \theta_{-j}) \left(\psi_j + \frac{\psi_j}{\psi_{-j}}\right) \ln H_{-j,t-i}}_{\text{Cross-Gender Work Norm}} \quad (\text{F.1})$$

The core empirical objective of this paper narrows down to testing a single hypothesis on the work norm for women:

$$H_0 : \theta_f = 0 \quad \text{vs.} \quad H_1 : \theta_f > 0 \quad (\text{F.2})$$

Under the null hypothesis ($H_0 : \theta_f = 0$), preferences lack a cultural transmission channel, meaning the historic rise in female labor force participation must be entirely explained by transitory, contemporary price shocks. Under the alternative hypothesis ($H_1 : \theta_f > 0$), female labor supply is governed by a self-reinforcing, cohort-specific cultural

engine. To test this mechanism, we leverage U.S. Census microdata and estimate the female participation equation directly, exploiting cross-state variation within specific cohorts.

Appendix F.2. Data and Empirical Sample

The empirical analysis is conducted using microdata from the Integrated Public Use Microdata Series (IPUMS USA; Ruggles et al., 2025). The core estimation sample is organized around two overlapping horizons designed to test the preference-transmission mechanism. Our baseline historical sample covers the years 1940 to 1980 ($N \approx 1.23$ million married couples), capturing the peak decades of the structural labor reallocation of married women. We subsequently extend the sample to track the 1940 to 2019 horizon ($N \approx 2.13$ million married couples) to evaluate the stability of the preference-transmission signature across more recent generations.

To isolate stable household lifecycle tracks and clean early-life preference signatures, the sample is restricted to married couples where both spouses are native-born US citizens. The analysis covers individuals belonging to birth cohorts born between 1885 and 1955. This wide historical window provides high-powered cross-cohort variation, contrasting generations socialized prior to World War I with those socialized during the peak mid-century transformation of female work norms.

Variable Construction and Measurement.

- *Labor Force Participation:* The primary outcome variable is the wife’s adult labor force participation, constructed as a binary indicator equal to 1 if she is currently employed or actively seeking work during the census reference week, and 0 otherwise. Across our baseline sample, this evaluates to a mean participation rate of 0.36. Intensive-margin hours adjustments are tracked via weekly hours worked (UHRSWORK) and annual weeks worked (WKSWORK2).
- *Formative Work Norm (H_{jt-i}):* Following the geographic, cohort-varying framework of Fogli and Veldkamp (2011), the gender-specific formative work norm is measured as the prevalence of female employment in an individual’s *birth state* during their childhood upbringing (ages 0 to 15). This geographic structure isolates an aggregate, localized cohort experience rather than an idiosyncratic household shock.
- *Cross-Spouse Exposure:* For the cross-spouse designs pioneered by Fernández et al. (2004), husbands are mapped to their own childhood environments using their respective birth state and birth year. This isolates how the work norms of a husband’s upbringing independently shape his future wife’s adult participation decisions.
- *Epidemiological Subsample:* To isolate pure preference transmission from contemporaneous local labor market institutions or institutional rules, a clean subsample of approximately 77,500 second-generation immigrant women (US-born individuals with foreign-born parents) is constructed for the years 1940–1960. Following Fernández et al. (2004), these individuals are assigned the historical female labor force participation rates of their parents’ country of origin across 47 distinct countries, evaluating transmitted preferences independent of US local markets.
- *Household Resources and Controls:* Annual wage and salary income (INCWAGE) is deflated using the Consumer Price Index (CPI) to control for real resource tracks. To rule out bargaining-power alternatives, the data are augmented with localized marriage-market sex ratios and a predetermined spousal education-gap control.

While household-specific lifetime resource variations remain, our cross-state within-cohort design isolates the component of early childhood exposure that is orthogonal to these local macroeconomic shocks. Contemporaneous, state-level labor demand shifters, institutional policy variations, and regional wealth innovations are non-parametrically absorbed by residence-state-by-year fixed effects (as implemented in Table 3, Row 5). Thus, the structural parameter governing cultural transmission remains identified through the compound structural loading $\theta_f(\psi_f + 1) \ln H_{f,t-i}$, safely insulated from time-varying regional confounders.