NBER WORKING PAPER SERIES

THE EMPATHY CHANNEL IN FERTILITY

Sebastian Galiani
Raul A. Sosa

Working Paper 35021
http://www.nber.org/papers/w35021

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2026

We used generative artificial intelligence tools (Claude by Anthropic and ChatGPT by OpenAI) for coding assistance, literature verification, and editorial support. All analytical content and substantive conclusions are the authors' own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

The Empathy Channel in Fertility
Sebastian Galiani and Raul A. Sosa
NBER Working Paper No. 35021
April 2026
JEL No. J13

ABSTRACT

Being around babies makes people want babies. We formalize this observation as the empathy channel: exposure to infants in the social environment activates neurobiological mechanisms that increase the desire for parenthood. As children become scarcer, this affective stimulus weakens, further eroding the motivation to have children. We embed the mechanism in a two-group overlapping-generations quantity-quality model. The empathy channel generates a positive externality, since each birth raises others' desire for children, making the decentralized equilibrium inefficient. We characterize the optimal per-child subsidy and show that the first-order Pigouvian rate substantially overshoots the general-equilibrium optimum. The optimal targeting rule follows a Ramsey-like logic, directing the subsidy at the group with the most externality per fiscal dollar, not the group with the largest externality per child. The calibrated model suggests that the empathy channel can account for 3–33% of the fertility decline, with 13.4% at the baseline. At this baseline, the Pigouvian overshoot is 23–32% and the optimal subsidy raises welfare by 0.22% in consumption-equivalent terms.

Sebastian Galiani
University of Maryland, College Park
Department of Economics
and NBER
sgaliani@umd.edu

Raul A. Sosa
Universidad de San Andres
Department of Economics
rsosa@udesa.edu.ar

1 Introduction

Fertility rates have fallen dramatically across the world. Rising education costs, expanding labor market opportunities for women (Goldin, 2025b), and improved contraceptive access (Goldin and Katz, 2002) all contribute, and the macroeconomic consequences may be severe (Jones, 2022). We propose that an additional force is at work, one grounded in biology, that may compound these economic pressures. Exposure to infants in the social environment activates an affective response that increases the desire for parenthood. Where adults see more children, adults want more children. Where the young population shrinks, so does the desire for parenthood.

In the canonical quantity-quality framework (Becker, 1960; Becker and Lewis, 1973), parents optimize over the number and quality of their children, and a cost shock reduces fertility through straightforward budget adjustment. The preference for children, however, is taken as given. When it responds to the social environment, as we argue here, each generation's fertility shapes the demographic environment that the next generation faces. A cost-driven decline can therefore compound across generations, amplifying the initial shock beyond what costs alone would predict.

We call this the empathy channel. Evidence reviewed in Section 2 establishes that infant exposure activates the brain's reward circuitry (Glocker et al., 2009), triggers oxytocin release promoting caregiving desire (Feldman, 2017), and extends to fathers and other non-birthing caregivers (Abraham et al., 2014). The mechanism operates through sensory contact rather than deliberation. Psychologists find that the affective component of fertility motivation is more than twice as strong as the cognitive one in predicting "baby fever" intensity (Brase and Brase, 2012).

We formalize this insight within an overlapping-generations quantity-quality model in the tradition of de la Croix and Doepke (2003). Two groups coexist, distinguished by their degree of exposure to infants. Preferences follow the standard Becker structure, augmented by a single additive term in the spirit of Brock and Durlauf (2001) that links the desire for children to the stock of children in the surrounding community. The empathy function is concave and saturating, motivated by the dose-dependent but

bounded character of the neurobiological response. When the empathy channel is shut down, the model reduces to a standard two-group quantity-quality framework.

The analysis delivers four results.

First, the empathy channel generates a positive externality and the decentralized equilibrium features subfertility. Each birth adds an infant to the social environment, raising the empathy signal for surrounding adults. No family internalizes this effect, making births a public good that the market underprovides. Our contribution is to provide a biological microfoundation for this social channel. Existing models of fertility contagion document that others' fertility decisions affect one's own and attribute the interdependence to social interactions. The biology we review identifies a specific pathway through which these interactions operate and distinguishes it from cognitive channels such as social learning and conformity.

Second, the first-order Pigouvian correction substantially overshoots the general-equilibrium (GE) optimum. The subsidy itself raises fertility, which raises exposure and empathy, partially closing the gap the subsidy was designed to correct. At the Pigouvian rates, the subsidy actually reduces welfare. The GE-optimal subsidy is substantially lower, and we characterize the marginal cost of public funds (MCPF) threshold below which it is warranted. The externality allows a baby bonus to withstand 72% higher fiscal costs than an equivalent cash transfer before becoming welfare-negative.

Third, the optimal subsidy targets the high-exposure group, following a Ramsey-like targeting rule. The low-exposure group generates a larger externality per child (because its population share is larger), but the high-exposure group generates a larger externality per fiscal dollar (because its population is smaller and its children are less expensive to subsidize). The optimal target is the group where the welfare return per dollar of spending is highest, not the group with the largest per-child externality. A uniform subsidy, however, captures over 98% of the group-specific welfare gain, providing a practical recommendation that avoids differentiated transfers.

Fourth, the calibrated model suggests that the empathy channel can account for 3–33% of the fertility decline, depending on the strength of the channel. At the baseline calibration, this

share is roughly 13%. An increase in child-rearing costs that would reduce aggregate fertility by 17.4% in the absence of empathy produces a 20.1% decline when the empathy channel is active. The range 3–33% reflects variation in the empathy parameter across values consistent with the peer-effects literature (Lyngstad and Prskawetz, 2010; Balbo and Barban, 2014; Büyükkeçeci et al., 2020). At this baseline, the optimal subsidy raises welfare by 0.22% in consumption-equivalent (CE) terms, approximately $117 per family per year, and survives distortionary financing at realistic MCPF levels. During the transition to the new steady state, the optimal subsidy is approximately constant at 40% of the long-run optimum, and allowing it to vary across generations adds no welfare gain.

Related literature. This paper connects four strands. The first is the empirical literature on fertility contagion and peer effects. Lyngstad and Prskawetz (2010) find that a sibling's birth more than doubles the hazard of first birth in Norwegian registry data. Balbo and Barban (2014) document significant fertility contagion among friends using U.S. panel data, with the effect peaking approximately two years after a friend's birth. Büyükkeçeci et al. (2020) document significant coworker effects in Dutch administrative data. Kuziemko (2006) shows that fertility contagion between siblings is strongest among the childless, consistent with the exposure mechanism we formalize. At the macro level, Kohler et al. (2002) argue that social interaction effects amplified the fertility decline to lowest-low levels in Southern and Eastern Europe, producing a social multiplier on modest socioeconomic changes. These papers establish that fertility decisions are socially interdependent and propose various channels, including social learning, conformity, and cost-sharing. Our model identifies a distinct biological mechanism, grounded in the neuroscience of infant exposure, and derives its welfare implications.

The second strand is the theoretical literature on fertility, human capital, and growth. Our model builds on the quantity-quality tradition initiated by Becker (1960) and Becker and Lewis (1973), in which parents trade off the number and quality of their

children. Galor and Weil (2000) embedded this tradeoff in a growth model to explain the demographic transition, and de la Croix and Doepke (2003) introduced group heterogeneity in child-rearing costs. We add a social utility term in the spirit of Brock and Durlauf (2001). The model complements household bargaining approaches (Doepke and Kindermann, 2019) and the recent survey of the economics of fertility by Doepke et al. (2023).

The third strand connects social interactions and culture to fertility. Manski (1993) formalized the identification challenge in distinguishing endogenous social effects from contextual and correlated effects. Brock and Durlauf (2001) provided the social interactions framework we build on, and Glaeser et al. (2003) formalized the social multiplier concept. Manski and Mayshar (2003) applied social interactions to fertility across religious groups in Israel. The culture literature demonstrates that fertility has a socially transmitted component operating through the social environment: Fernandez and Fogli (2009) show that ancestry-country fertility predicts second-generation women's childbearing, with the effect amplified by co-ethnic density. Fernandez and Fogli (2006) distinguish culture from direct family experience, finding that both matter independently. Bisin and Verdier (2000, 2001) model the intergenerational transmission of cultural traits through endogenous socialization. Galiani and Sosa (2025) extend this framework to study how differential fertility and cultural retention drive long-run compositional dynamics. Spolaore and Wacziarg (2022) document that the European fertility transition diffused along cultural-proximity lines, not along industrialization lines. Lesthaeghe (2010) frames the broader shift as a Second Demographic Transition driven by value change. The empathy channel adds a biological pathway through which declining fertility itself erodes the desire for parenthood, a mechanism absent from informational and cultural accounts.

A fourth strand addresses the normative question of what policy should do about declining fertility. Kearney and Levine (2025) survey the contemporary decline and frame it as a demand-side puzzle, calling for research into the determinants of fertility preferences. Gauthier (2007) finds that cash transfers have had modest effects on

fertility, consistent with the view that the decline reflects shifts in preferences rather than constraints alone. Goldin (2025a) argues that countries experiencing rapid post-war economic growth with lagging cultural adaptation became lowest-low fertility nations, a pattern consistent with the mechanism we formalize. The welfare analysis grounds the case for intervention in externality correction rather than demographic targets.

The rest of the paper is organized as follows. Section 2 reviews the biological evidence. Section 3 presents the model. Section 4 develops the welfare framework. Section 5 calibrates and quantifies. Section 6 concludes.

2 The Biology of Fertility Motivation

Why did evolution produce not only sexual desire but also love for children? Reproduction, in principle, requires only libido. Many mammals reproduce successfully without the extended, intensive parental investment that human children require. But human infants are extraordinarily dependent: large-brained, slow to mature, and unable to survive without sustained care for years. A species with this life history cannot rely on sexual drive alone to ensure the survival of its young. Selection pressures therefore favored mechanisms that generate not merely reproduction but *commitment*: a sustained motivation to invest in children that goes beyond the act of conception. Understanding these mechanisms is not optional for a general theory of fertility.

This section reviews five bodies of evidence, from neuroscience, endocrinology, psychology, evolutionary biology, and anthropology, that converge on a single conclusion. A measurable component of fertility motivation is activated by exposure to infants in the social environment. The mechanism is affective, not cognitive. It operates through sensory contact rather than information or deliberation, and it is social: the relevant stimulus extends beyond one's own children to those present in one's community.

Neuroscience of infant perception. Lorenz (1943) first identified the *Kindchenschema*, or baby schema: a constellation of infantile facial features that functions as an innate

releasing mechanism, automatically triggering caregiving behavior without requiring learning or prior experience. The response extends to juvenile animals and dolls, suggesting deep evolutionary origins. Glocker et al. (2009) provided the first neural evidence using functional magnetic resonance imaging (fMRI). Photographs of infant faces with high baby schema intensity produced significantly greater activation in the nucleus accumbens, a core structure of the brain's reward circuit, in women who had never given birth. The reward response is thus part of the species-typical neural architecture, not something acquired through motherhood, and is graded: more intense baby features produce stronger responses, consistent with a dose-dependent mechanism. Kringelbach et al. (2016) showed that this response operates through rapid attentional capture within approximately 140 milliseconds, before conscious processing, followed by sustained reward activation. When an adult sees a baby, the brain's reward system engages before conscious deliberation begins.

The oxytocin system. The reward response to infant features is the entry point, but the deeper pathway runs through the oxytocin system. Oxytocin is a neuropeptide that plays a central role in social bonding, attachment, and caregiving. Ruth Feldman's programmatic research (Feldman, 2012, 2017) has established that plasma oxytocin levels correlate with the degree of infant contact, that oxytocin release promotes bonding through a bidirectional loop (contact promotes oxytocin, which promotes bonding, which generates more contact), and that the system extends beyond the mother-infant dyad to paternal and alloparental bonds (Hrdy, 2009). Abraham et al. (2014) confirmed that caregiving is not confined to biological mothers: primary-caregiving mothers, secondary-caregiving fathers, and primary-caregiving fathers in same-sex couples all showed activation of a parental caregiving network, with functional connectivity in fathers modulated by caregiving experience. Hands-on experience with infants is sufficient to activate the same neural infrastructure regardless of biological parenthood.

Empathic simulation. A further body of evidence concerns empathic simulation: the ability to internally represent another person's emotional state through obser-

vation alone. Singer and Lamm (2009) document that observing another person's emotional state activates shared neural representations in the observer, though the precise mechanisms remain debated (Hickok, 2009). If observing parental behavior activates caregiving-related affect, then the channel extends beyond direct infant contact to vicarious experience: seeing others care for children partially engages the same affective circuits. This is not social learning in the informational sense but affective contagion, and the distinction matters because information can be transmitted through statistics and conversations, while the affective response requires sensory proximity to the stimulus.

Baby fever and fertility motivation. The evidence reviewed so far establishes that infant exposure activates caregiving affect. The psychological literature extends this finding to fertility motivation specifically. Rotkirch (2007) introduced the concept of "baby longing": a visceral, embodied desire for a child experienced as a physical sensation rather than as the outcome of deliberation. The desire is surprising, involuntary, and sometimes at odds with rational cost-benefit assessment. Rotkirch et al. (2011) documented similar desires in men, in some cases triggered by exposure to friends' or siblings' children, a finding that cannot be attributed to pregnancy-related hormonal changes. Brase and Brase (2012) conducted the most systematic investigation, decomposing fertility motivation into three independent components: a positive exposure factor (emotional responses to seeing babies), a negative exposure factor (aversion to childcare), and a tradeoffs factor (deliberative cost-benefit assessment). The positive exposure component was more than twice as strong as the tradeoffs component in predicting baby fever intensity, suggesting that fertility motivation has a visceral, pre-deliberative component that operates alongside the weighing of costs and benefits.

Hrdy (1999), in *Mother Nature*, demonstrates that maternal commitment is not a binary switch but a graded response, conditional on the social environment. Across primate species and human history, maternal investment has varied with the availability of social support. In traditional societies, girls and young women arrive at their first

pregnancy having accumulated extensive hands-on experience with infants through caring for younger siblings and neighbors' children, an apprenticeship that prepares the neurobiological systems for motherhood. The desire for children, on this account, is cultivated through experience.

Cooperative breeding. The evolutionary context explains why these mechanisms exist. Hrdy (2009) argues that humans are *cooperative breeders*: unlike other great apes, human mothers depend on alloparental assistance from fathers, grandmothers, older siblings, and unrelated community members to rear their extraordinarily dependent offspring. The extended juvenile dependence of human children, lasting fifteen to twenty years, requires more sustained investment than a single mother can provide. In the ancestral environment, adults routinely cared for other people's children, keeping infant exposure high and widely distributed across the community. Sear and Mace (2008), reviewing 45 studies across natural fertility populations, confirm that the presence of alloparents is consistently associated with improved child survival. Burkart et al. (2009) synthesize comparative evidence showing that cooperative breeders among primates are significantly more prosocial than non-cooperative breeders, consistent with selection for voluntary care of others' offspring. Kramer (2010) connects cooperative breeding to the demographic success of the species: the cooperative system allowed humans to achieve shorter interbirth intervals than other great apes while maintaining the extended juvenile period.

The modern departure. Modern industrial societies represent a sharp departure from this ancestral pattern. Fertility itself has declined, meaning fewer babies in the environment. Residential patterns have become age-segregated, childcare has been privatized within the home or specialized facilities, and the apprenticeship in infant care documented by Hrdy (1999) has largely disappeared. The cross-cultural evidence is consistent: more religiously observant populations, characterized by stronger community ties and higher infant visibility, maintain higher fertility (Frejka and Westoff, 2008). The Amish provide a particularly clear illustration: concentrated settlement patterns,

extensive within-community childcare, and high infant visibility are associated with persistently high fertility (Ericksen et al., 1979; Greksa, 2002; Wasao et al., 2021), despite access to modern contraception.

Taken together, these five bodies of evidence point to a conclusion with direct implications for fertility theory. A component of the motivation to have children is activated by exposure to infants in the social environment. The mechanism is biological, grounded in reward circuitry and the oxytocin system. It is social, because the relevant stimulus is the stock of children in the community. And it is endogenous, because exposure depends on others' fertility decisions. Fertility does not simply respond to economic fundamentals. It evolves through a feedback loop in which the social environment, shaped by past reproductive decisions, influences the motivation for future ones.

3 Model

This section formalizes the biological mechanism documented in Section 2 within a quantity-quality fertility model. We extend the standard quantity-quality preference structure (Becker, 1960; Becker and Lewis, 1973) by a single additive term in which the desire for children depends on the stock of children in the agent's social environment. When the empathy parameter is set to zero, the model reduces to a standard two-group quantity-quality framework.

3.1 Environment

The economy is populated by overlapping generations (OLG). Each individual lives for two periods: a childhood period in which the individual is passive and an adulthood period in which economic decisions are made. Two groups coexist, denoted H (high-exposure) and L (low-exposure), with population shares λ_H and $\lambda_L = 1 - \lambda_H$. Group H represents communities with concentrated residential patterns and strong within-group social contact, while Group L represents populations with more dispersed social

networks.

In each period, adults observe the children present in their social environment, experience an affective response, and make fertility and investment decisions. Because the demographic environment was shaped by previous generations' fertility, an intergenerational feedback loop emerges: current fertility shapes the environment that will influence future fertility decisions.

Each adult is endowed with one unit of income ($w = 1$) and allocates it across consumption c_i, a base child-rearing cost ϕ_i per child, and a quality investment γ per unit of quality per child:

$$c_i = w - \phi_i \cdot n_i - \gamma \cdot n_i \cdot h_i, \tag{1}$$

where n_i is the number of children, h_i is quality investment per child, $\phi_i > 0$ is the group-specific per-child cost, and $\gamma > 0$ is the per-unit quality cost (common across groups). The base cost ϕ_i is lower for Group H ($\phi_H < \phi_L$), reflecting communal child-rearing and lower opportunity cost in more cohesive communities. Since we normalize income to unity for both groups, group-specific costs capture this heterogeneity in reduced form.

3.2 Preferences and the Empathy Function

An adult in group $i \in \{H, L\}$ derives utility from three sources: consumption, a composite of child quantity and quality, and an empathy-driven bonus to the value of children. The utility function is

$$U_i = \underbrace{\ln(c_i)}_{\text{consumption}} + \underbrace{\theta\left[\nu_i \ln(n_i) + (1-\nu_i)\ln(h_i)\right]}_{\text{quantity-quality composite}} + \underbrace{\psi \cdot e_i \cdot \ln(n_i)}_{\text{empathy channel}}. \tag{2}$$

The first two components are standard. The parameter $\theta > 0$ governs the overall weight on children (common across groups, following de la Croix and Doepke, 2003), and $\nu_i \in (0, 1)$ governs the quantity-quality split within each group. The third component is the empathy channel: an additional pronatal motive whose strength depends on the

empathy signal $e_i \in [0, 1)$. This term enters additively, following the general principle in the social interactions literature (Brock and Durlauf, 2001) that individual payoffs may include a component that depends on aggregate behavior. When $\psi = 0$, the empathy channel shuts down and the model reduces to a standard two-group quantity-quality model.

The total weight on $\ln(n_i)$ in utility is therefore $\theta\nu_i + \psi e_i$, the sum of intrinsic preference and empathy. We define the **net quantity premium**

$$A_i \equiv \theta(2\nu_i - 1) + \psi\, e_i, \tag{3}$$

which measures the excess of the effective quantity weight over the quality weight $\theta(1 - \nu_i)$. This composite parameter plays a central role in the closed-form solution: it determines whether fertility is positive, the direction of comparative statics, and the magnitude of the externality.

The empathy function. The empathy signal e_i is a concave, bounded function of effective exposure E_i:

$$e(E) = 1 - \exp(-\rho \cdot E), \qquad e : \mathbb{R}_+ \to [0, 1), \tag{4}$$

where $\rho > 0$ is the single parameter governing the rate at which exposure translates into empathy. The function has five properties that make it well suited to the biological evidence. It starts at zero: absent exposure, there is no empathy signal. It is strictly increasing: more babies in the environment produce a stronger affective response. It is strictly concave: the first babies an adult encounters have the strongest effect, with additional exposure yielding diminishing returns, consistent with the dose-dependent but saturating character of oxytocin release (Feldman, 2017). It is bounded above by unity: the empathy motive can amplify but never dominate preferences. Finally, its initial slope equals ρ, which controls the marginal sensitivity at zero exposure.

The concavity of the empathy function is the source of the model's most distinctive

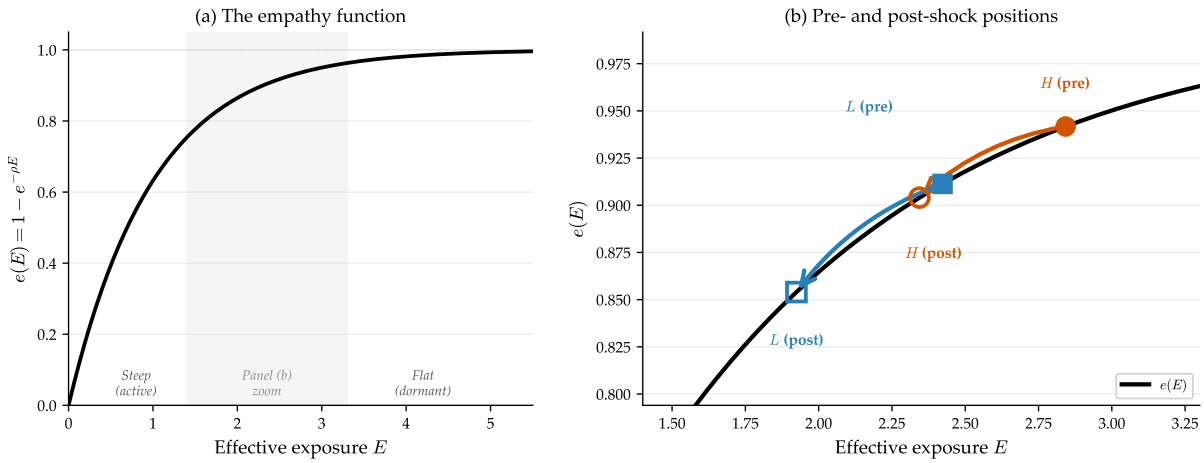

Figure 1. *The empathy function $e(E) = 1 - \exp(-\rho E)$ with $\rho = 1$. Panel (a): the full curve, concave and bounded above by 1. Panel (b): zoom on the operating range, showing how both groups slide leftward from their near-saturated pre-shock positions ($e > 0.91$) into the steeper region after the cost shock, where the marginal sensitivity $e'(E)$ is larger and the channel activates.*

dynamic prediction. At high fertility levels, such as those prevailing at the pre-shock steady state, the empathy signal is near saturation ($e > 0.91$ at calibrated parameters). The marginal response $e'(E) = \rho \exp(-\rho E)$ is small, and the channel is effectively dormant: a marginal birth barely registers in the empathy of surrounding adults. As fertility declines and exposure falls, the operating point moves leftward along the curve into the steeper region, and the marginal sensitivity rises sharply. The empathy channel, dormant at high fertility, activates during the decline (Figure 1). This asymmetry distinguishes the empathy mechanism from conformity models, which predict symmetric effects in both directions.

3.3 Exposure

The empathy function translates exposure into affect, but exposure itself depends on the demographic composition of an agent's social environment. We model this as a weighted combination of the local (within-group) child density and the aggregate child density:

$$E_i = \delta_i \cdot N_i + (1 - \delta_i) \cdot \bar{N}, \qquad \bar{N} \equiv \lambda_H N_H + \lambda_L N_L, \tag{5}$$

where N_i is the density of children in group i's environment (taken as given by each agent), $\bar{N}$ is the population-weighted aggregate child density, and $\delta_i \in [0,1]$ is the segregation parameter. The parameter δ_i governs how much of an agent's exposure comes from their own community versus the broader population. We assume $\delta_H > \delta_L$: high-exposure communities are more insular, consistent with the concentrated residential patterns and strong community ties that characterize the populations Group H represents.

The structure of exposure has direct consequences for how a birth in one group affects the other. The marginal effect of group i's child stock on group j's exposure is

$$\frac{\partial E_j}{\partial N_i} = \delta_j \cdot \mathbf{1}_{j=i} + (1 - \delta_j)\lambda_i. \tag{6}$$

This has two components: a within-group channel (active only when $j = i$, with weight δ_j) and a cross-group channel (always active, with weight $(1 - \delta_j)\lambda_i$). A birth in Group H raises H's own exposure by a large amount (through the within-group channel) and raises L's exposure by a smaller amount (through the cross-group channel, proportional to H's population share λ_H). These marginal effects enter the externality analysis in Section 3.7.

3.4 Optimal Choices

Each agent takes the empathy signal e_i as given and chooses n_i and h_i to maximize (2) subject to (1). This is a static optimization problem, which we solve here and derive in Appendix A.

The first-order condition (FOC) for quality yields a familiar Cobb-Douglas property: total quality spending is a fixed fraction of consumption,

$$\gamma\, n_i\, h_i = \theta(1 - \nu_i) \cdot c_i. \tag{7}$$

That is, the household devotes a constant share of its consumption to quality investment. The intuition is that with log preferences over quality, the marginal utility of quality

spending relative to consumption is always equalized at the same ratio, independent of costs.

The first-order condition for quantity sets the marginal benefit of an additional child (intrinsic preference plus empathy) equal to its marginal cost relative to consumption. Combining both conditions with the budget constraint and collecting terms yields the closed-form solution. Define $D_i \equiv A_i + 1 + \theta(1 - \nu_i)$. Substituting the definition of A_i, this simplifies to $D_i = \theta\nu_i + 1 + \psi e_i$, the "total expenditure coefficient." Then:

$$n_i^* = \frac{A_i \cdot w}{\phi_i D_i}, \qquad h_i^* = \frac{\theta(1-\nu_i) \cdot \phi_i}{\gamma \cdot A_i}, \qquad c_i^* = \frac{w}{D_i}. \tag{8}$$

Optimal fertility n_i^* is the net quantity premium A_i (how much the agent values children above quality) scaled by income and divided by the per-child cost, with the expenditure coefficient D_i capturing the crowding-out from all uses of income. Optimal quality h_i^* is proportional to the quality preference weight $\theta(1 - \nu_i)$ and inversely proportional to A_i: agents who value quantity more invest less per child, the classic quantity-quality tradeoff. Optimal consumption c_i^* is income divided by D_i, the share left over after child-related spending.

At the optimum, income divides into three components in the ratio $1 : A_i : \theta(1 - \nu_i)$, corresponding to consumption, base child-rearing, and quality investment. This ratio depends only on preference parameters and the empathy signal, not on costs or income.

3.5 Properties

Proposition 3.1 (Interior Optimum)**.** *If $A_i > 0$, $\nu_i \in (0, 1)$, $\phi_i > 0$, $\gamma > 0$, and $w > 0$, the optimal choices satisfy $n_i^* > 0$, $h_i^* > 0$, and $c_i^* > 0$. The Hessian of the reduced-form utility is negative definite at this interior solution.*

The proof is in Appendix A. The result is immediate from inspection of (8): all three expressions are strictly positive when $A_i > 0$. The second-order conditions confirm that the solution is a maximum, not a saddle point.

The role of A_i. The condition $A_i > 0$ deserves scrutiny because it determines whether the model produces positive fertility. Without empathy ($\psi = 0$), the condition requires $\nu_i > \frac{1}{2}$: the intrinsic quantity weight must exceed the quality weight. With empathy ($\psi > 0$), the condition relaxes to $\nu_i > \frac{1}{2} - \psi e_i/(2\theta)$. Empathy can sustain positive fertility even for quality-oriented agents with $\nu_i < \frac{1}{2}$, because the empathy signal provides an additional pro-quantity motive beyond intrinsic preferences. At calibrated parameters, both groups satisfy $\nu_i > \frac{1}{2}$, so $A_i > 0$ holds for all exposure levels, guaranteeing a well-defined interior solution everywhere in the state space.

Independence from quality costs. A feature of the Cobb-Douglas structure is that optimal fertility n_i^* does not depend on the quality cost parameter γ. This is immediate from (8): γ appears only in h_i^*, not in n_i^*. Intuitively, the budget share allocated to quality is fixed by preferences. Changing γ rescales quality units (the "amount" of quality per dollar) but not quality spending, leaving fertility unaffected. This separation allows γ to be calibrated independently of fertility targets.

Comparative statics. Higher child costs reduce fertility with unit elasticity and raise quality. Higher empathy raises fertility and reduces quality, reflecting the quantity-quality tradeoff.

3.6 Dynamics

The OLG structure generates dynamics through the intergenerational feedback loop: current fertility determines the child stock that will shape the next generation's empathy and fertility decisions. Denoting the state vector $\mathbf{N}_t = (N_{H,t}, N_{L,t})$, the dynamical system is

$$\mathbf{N}_{t+1} = F(\mathbf{N}_t), \qquad F_i(\mathbf{N}) = n_i^*\big(e(E_i(\mathbf{N}))\big). \tag{9}$$

The map F defines a two-dimensional system: the child stock in each group next period equals the fertility choice of the current generation, which depends on empathy, which depends on the current child stock. The chain of causation runs from child stock to

exposure to empathy to fertility to next period's child stock.

A steady state (N_H^*, N_L^*) is a fixed point of this map: a pair of child densities such that, given the empathy signals they generate, each group's optimal fertility reproduces exactly its own child density.

Proposition 3.2 (Existence and Uniqueness)**.** *If $v_i > \frac{1}{2}$ for both groups, at least one steady state exists with $N_i^* > 0$ for both i. At calibrated parameters, the steady state is unique.*

Existence follows from a standard fixed-point argument. Under $v_i > \frac{1}{2}$, the condition $A_i > 0$ holds for all exposure levels, so $F_i(\mathbf{N}) > 0$ everywhere, including at zero. The map F is continuous and sends the compact rectangle $[0, w/\phi_H] \times [0, w/\phi_L]$ into its interior. By Brouwer's fixed-point theorem, a fixed point exists. Uniqueness is verified numerically at calibrated parameters via grid search (see Appendix A). The concave empathy function, which saturates quickly, prevents the multiple equilibria that can arise in social interaction models with S-shaped conformity functions (cf. the low-fertility trap of Lutz et al., 2006).

Remark 3.3 (Stability)**.** *Local asymptotic stability requires both eigenvalues of the Jacobian $J = D_{\mathbf{N}}F$ to have modulus less than one. At calibrated parameters, the spectral radius is 0.123, confirming fast convergence.*

The Jacobian has entries

$$J_{ij} = \frac{\partial n_i^*}{\partial N_j} = \underbrace{\frac{\partial n_i^*}{\partial A_i}}_{\substack{\text{fertility} \\ \text{response}}} \cdot \underbrace{\psi \cdot \rho \, \exp(-\rho E_i)}_{\substack{\text{empathy} \\ \text{sensitivity}}} \cdot \underbrace{\frac{\partial E_i}{\partial N_j}}_{\substack{\text{exposure} \\ \text{structure}}} . \tag{10}$$

Each entry decomposes into three factors. The fertility response captures how strongly optimal n_i responds to a change in A_i. The empathy sensitivity captures how responsive the empathy function is at the current exposure level, namely the marginal $e'(E_i) = \rho \exp(-\rho E_i)$, weighted by the channel strength ψ. The exposure structure captures the network through which a change in group j's child stock reaches group i's exposure. All entries are non-negative, confirming the cooperative structure: higher fertility in either group raises empathy and fertility in both groups.

At calibrated parameters, the empathy function is near saturation at the steady state, so the empathy sensitivity factor is small ($e'(E) \approx 0.06$). This keeps the spectral radius of J well below unity (0.123 at the baseline), implying fast convergence: the system reaches the new steady state within two to three generations after a perturbation. The stability reflects a fundamental feature of the model: the concave empathy function creates diminishing, not increasing, returns to fertility, so the feedback loop is self-limiting rather than explosive.

The social multiplier. The total effect of a cost shock on steady-state fertility exceeds the direct effect due to the empathy feedback. Totally differentiating the steady-state system:

$$d\mathbf{N} = \underbrace{(\mathbf{I} - J)^{-1}}_{\text{social multiplier}} \cdot \underbrace{D_{\boldsymbol{\phi}} F}_{\text{direct effect}} \cdot d\boldsymbol{\phi}. \tag{11}$$

The matrix $(\mathbf{I} - J)^{-1}$ amplifies the direct cost effect: the amplification factor is the social multiplier (Glaeser et al., 2003). At calibrated parameters, the social multiplier is approximately 1.10–1.12×, conservative relative to the substantial dyadic effects documented in the peer-effects literature, where a sibling's or friend's birth can more than double the short-run hazard of first parenthood (Lyngstad and Prskawetz, 2010; Balbo and Barban, 2014) (though dyadic hazard ratios and steady-state multipliers are not directly comparable).

3.7 The Externality

The model generates a positive externality because each birth increases the child stock visible to other adults, raising their empathy and desire for children. Individual agents take the empathy signal as given and do not account for this effect. The decentralized equilibrium is therefore inefficient: it features lower fertility than the allocation a social planner would choose.

The private decision. In the decentralized equilibrium, each agent's first-order condition for n_i is

$$\frac{\theta v_i + \psi e_i}{n_i} = \frac{\phi_i + \gamma h_i}{c_i}. \tag{12}$$

The agent recognizes the empathy signal but treats it as exogenous. The crucial gap is that the agent does not internalize the effect of their own fertility on the child stock, and hence on the empathy signal and utility of other agents.

The social planner. Consider a planner who maximizes steady-state social welfare, recognizing that in the steady state the child stock equals current fertility ($N_i = n_i$). The planner chooses n_H, n_L, h_H, h_L to maximize

$$W = \sum_{i \in \{H,L\}} \lambda_i \left[\ln(c_i) + \theta \left[v_i \ln(n_i) + (1 - v_i) \ln(h_i) \right] + \psi \cdot e_i(n_H, n_L) \cdot \ln(n_i) \right]. \tag{13}$$

The key difference from the private problem is that the planner recognizes the dependence of e_j on n_i for all j. Differentiating with respect to n_i yields

$$\frac{\partial W}{\partial n_i} = \lambda_i \cdot \underbrace{\left[\frac{\theta v_i + \psi e_i}{n_i} - \frac{\phi_i + \gamma h_i}{c_i} \right]}_{\text{private net benefit (= 0 in decentralized)}} + \underbrace{\sum_j \lambda_j \cdot \psi \ln(n_j) \cdot \frac{\partial e_j}{\partial n_i}}_{\text{marginal external benefit}}. \tag{14}$$

The first term is the private marginal net benefit, which equals zero in the decentralized equilibrium. The second term is the **marginal external benefit** (MEB) of a child from group i: the welfare effect of n_i on everyone's empathy.

Expanding the MEB gives the key equation:

$$\text{MEB}_i = \sum_{j \in \{H,L\}} \lambda_j \cdot \underbrace{\psi \ln(n_j)}_{\text{empathy value}} \cdot \underbrace{\rho \exp(-\rho E_j)}_{\text{empathy sensitivity}} \cdot \underbrace{\frac{\partial E_j}{\partial N_i}}_{\text{exposure propagation}}. \tag{15}$$

Three factors determine the external benefit of a birth. The *empathy value* $\psi \ln(n_j)$ captures how much agent j's utility responds to a change in empathy, and it is positive whenever $n_j > 1$. The *empathy sensitivity* $e'(E_j) = \rho \exp(-\rho E_j)$ captures how responsive

the empathy function is at group j's current exposure level. It is large when exposure is low (the steep region of the concave function) and small when exposure is high. The *exposure propagation* $\partial E_j / \partial N_i$ captures the network structure: how much group i's child stock contributes to group j's exposure, through both within-group and cross-group channels.

Proposition 3.4 (Positive Externality and Subfertility)**.** *When $\psi > 0$ and $n_i^* > 1$ at the steady state,* $\text{MEB}_i > 0$ *for both groups. The decentralized equilibrium features lower fertility than the first-best allocation.*

The proof is immediate from the decomposition in (15): every factor is strictly positive. The empathy value is positive because $\psi > 0$ and $\ln(n_j) > 0$ (fertility exceeds one at calibrated parameters). The empathy sensitivity is positive for any finite exposure. The exposure propagation is positive because both the within-group and cross-group channels transmit positive effects. Since $\text{MEB}_i > 0$ and the private equilibrium sets the private net benefit to zero, the social marginal benefit exceeds the private marginal benefit at the decentralized allocation. A planner would choose higher fertility.

The direction of the externality is unambiguous: births are a public good. Each child adds to the stock of infants visible in the social environment, generating empathy that raises the utility of other adults. This benefit is not internalized in private fertility decisions. The result provides a justification for pronatal policy, one that does not require appeals to population targets or national interest but follows directly from the structure of preferences and the failure of agents to internalize the social consequences of their choices.

Only the quantity margin is distorted. The externality operates exclusively through the quantity channel: the planner's FOC for quality h_i is identical to the private condition (7), because empathy depends on the child stock, not on quality investment. A per-child subsidy is therefore a sufficient corrective instrument: the allocation induced by the optimal subsidy coincides with the first-best.

The Pigouvian correction. A per-child subsidy τ_i that reduces the effective cost from ϕ_i to $\phi_i - \tau_i$ yields the first-order Pigouvian rate:

$$\tau_i^{\text{Pigou}} = \frac{c_i \cdot \text{MEB}_i}{\lambda_i}, \tag{16}$$

evaluated at the decentralized allocation. This is a useful benchmark but only a first-order approximation. The full optimal subsidy, accounting for GE feedback and distortionary financing, is characterized in Section 4.

The targeting principle. The MEB differs across groups, creating a natural ranking for policy:

Definition 3.5 (Externality Rank)**.** *Group i has a higher externality rank than group j if* $\text{MEB}_i / (\lambda_i \phi_i) > \text{MEB}_j / (\lambda_j \phi_j)$*, that is, if an additional dollar of subsidy directed at group i generates more total external benefit.*

This ranking need not coincide with the ranking of MEB per child. The distinction is the basis of the targeting result developed in Section 4.

The externality identified above calls for corrective policy.

4 Welfare and Optimal Policy

Section 3.7 established that the decentralized equilibrium features subfertility and derived the first-order Pigouvian correction. But how should a government implement pronatal policy when raising revenue is costly, when groups differ in the return to intervention, and when the externality itself responds to the policy? This section develops the welfare framework. We characterize the MCPF threshold, the targeting rule, the welfare loss from uniformity, and the Pigouvian overshoot. The quantitative evaluation is deferred to Section 5.

4.1 The Planner's Problem

Objective. The planner maximizes steady-state utilitarian welfare:

$$W = \sum_{i \in \{H,L\}} \lambda_i U_i, \tag{17}$$

recognizing that in the steady state the child stock equals current fertility ($N_i = n_i$), so the empathy signal is endogenous to the allocation. The planner cannot directly mandate fertility or quality choices. Instead, the planner chooses per-child subsidies $\boldsymbol{\tau} = (\tau_H, \tau_L)$ that modify agents' budget constraints, and agents optimize privately given the subsidy. The planner uses price instruments, not quantity mandates.

Remark 4.1 (Empathy in the welfare function)**.** *The planner evaluates welfare using agents' empathy-modified preferences, including the $\psi e_i \ln(n_i)$ term. This is a substantive normative choice that deserves explicit defense. Three arguments support it.*

First, the empathy response is an evolutionary adaptation for cooperative breeding (Hrdy, 2009), part of the species-typical neural architecture, not a cognitive bias. It is no more a "distortion" than the taste for food or for social connection.

Second, in the framework of Bernheim and Rangel (2009), welfare-relevant preferences are those that survive deliberation and do not generate regret. Unlike addiction, where agents regret consumption once the visceral state passes, parents rarely regret children whose conception was motivated by empathy.

Third, on welfarist grounds (Stigler and Becker, 1977: de gustibus), if empathy enters the utility function, the planner respects it. The alternative, overriding empathy-modified preferences, requires a paternalistic judgment about which components of utility are "real."

Of course, this normative choice is not uncontested. One could instead classify empathy-induced desires as bias rather than genuine preference, which would alter the welfare conclusions. However, such a position carries its own burden: it requires a principled criterion for separating "authentic" from "influenced" components of utility. As Stigler and Becker (1977) argued, welfare economics has good reasons to avoid such a distinction. The biological evidence reviewed in Section 2 supports treating the empathy response as part of agents' stable preference

architecture.

Modified budget constraint. With a per-child subsidy τ_i financed by a lump-sum tax T, each agent's budget becomes

$$c_i = w - T - (\phi_i - \tau_i)\, n_i - \gamma\, n_i\, h_i. \tag{18}$$

The agent's problem is identical to the baseline but with effective parameters $\tilde{w} = w - T$ and $\tilde{\phi}_i = \phi_i - \tau_i$. The closed-form solutions become:

$$n_i^*(\tau_i, T) = \frac{A_i \cdot (w - T)}{(\phi_i - \tau_i) \cdot D_i} \tag{19}$$

$$h_i^*(\tau_i) = \frac{\theta(1 - \nu_i)(\phi_i - \tau_i)}{\gamma A_i} \tag{20}$$

$$c_i^*(\tau_i, T) = \frac{w - T}{D_i}. \tag{21}$$

Two features of these expressions deserve emphasis. First, the subsidy raises fertility and *lowers* quality. This is the quantity-quality tradeoff in action: by reducing the effective cost of children, the subsidy tilts the household toward more children with less investment in each. The quality reduction is not a first-order inefficiency (recall that the quality margin has no externality wedge), but it is a second-order distortion inherent in the instrument. Second, the empathy signal e_i (and hence A_i, D_i) depends on the steady-state child stock, which itself depends on the subsidy. The optimal subsidy is therefore characterized by a fixed-point condition, not a simple formula.

Government budget constraint. With lump-sum financing, the government balances its budget:

$$T = \sum_i \lambda_i \cdot \tau_i \cdot n_i^*(\tau_i, T). \tag{22}$$

The planner's problem is to choose (τ_H, τ_L) to maximize

$$\max_{\tau_H, \tau_L} \; W(\tau_H, \tau_L) = \sum_i \lambda_i\, U_i^*(\tau_i, T(\tau_H, \tau_L)), \tag{23}$$

where U_i^* is the indirect utility at the subsidized steady-state equilibrium, and T is implicitly defined by (22).

4.2 The Pigouvian Subsidy

First-order conditions. Differentiating W with respect to τ_i, three channels transmit the effect of a marginal subsidy increase to social welfare.

Direct channel. By the envelope theorem, a marginal increase in τ_i raises group i's consumption utility by n_i/c_i (one additional dollar per child, valued at the marginal utility of income).

Tax channel. The lump-sum tax rises to finance the subsidy, reducing everyone's consumption utility by $1/c_j$ per dollar of additional tax.

Externality channel. The subsidy changes n_i, which changes the steady-state child stock, which changes empathy and utility for all groups.

Combining these channels, the first-order condition at $\boldsymbol{\tau} = \mathbf{0}$ (the decentralized equilibrium) is

$$\left.\frac{\partial W}{\partial \tau_i}\right|_{\tau=0} = \lambda_i \frac{n_i}{c_i} - \lambda_i n_i \sum_j \frac{\lambda_j}{c_j} + \frac{n_i}{\phi_i} \cdot \mathrm{MEB}_i, \tag{24}$$

where the first two terms capture the distributional effect of a targeted transfer (private benefit minus tax cost) and the third term is the externality benefit, with $n_i/\phi_i = \partial n_i/\partial \tau_i|_{\tau=0}$ being the unit-elastic fertility response at zero subsidy.

The Pigouvian rate. To find the optimal subsidy, we compare the planner's first-order condition with the agent's. The planner internalizes the empathy externality, so the planner's optimal fertility satisfies an FOC with an additional term $\psi \ln(n_j)\, e'(E_j)\, (\partial E_j/\partial n_i)$ for each group j. The per-child subsidy that induces the agent to replicate the planner's

allocation must introduce a wedge τ_i/c_i equal to the externality benefit MEB_i/λ_i:

$$\frac{\tau_i}{c_i} = \frac{\text{MEB}_i}{\lambda_i}. \tag{25}$$

This is an implicit characterization: c_i and MEB_i depend on τ_i through the subsidized equilibrium. Evaluated at the decentralized allocation ($\boldsymbol{\tau} = \mathbf{0}$), it yields the first-order Pigouvian rate $\tau_i^{\text{Pigou}} = c_i \cdot \text{MEB}_i/\lambda_i$ already stated in (16).

Proposition 4.2 (Positive Subsidy)**.** *When $\psi > 0$ and $n_j > 1$ for all j at the steady state, the optimal per-child subsidy satisfies $\tau_i^* > 0$ for both groups.*

The proof follows immediately from $\text{MEB}_i > 0$ (Proposition 3.4), $c_i > 0$, and $\lambda_i > 0$: the right-hand side of (25) is strictly positive, so the subsidy must be as well. When $\psi = 0$, the externality vanishes and $\tau_i^* = 0$: the standard quantity-quality model requires no corrective policy.

Two remarks on the scope of this result. First, the condition $n_j > 1$ ensures $\ln(n_j) > 0$ so that the empathy term contributes positively to welfare. This is not a demographic replacement condition (which is approximately 2.1) but a mathematical condition for the sign of the externality. At calibrated parameters, both groups comfortably satisfy $n_j > 1$ even after the cost shock ($n_H = 2.50$, $n_L = 1.80$). Second, the subsidy is strictly increasing in ψ: a stronger empathy channel implies a larger externality and hence a larger correction. Although stronger empathy also reduces c_i (more income goes to children), the linear dependence of the MEB on ψ dominates the concave decrease in consumption. The formal argument is in Appendix A.

Why the Pigouvian rate overshoots. The first-order Pigouvian formula evaluates the externality at the *decentralized* allocation. But the subsidy raises fertility, which raises exposure and empathy, partially closing the gap that the subsidy was designed to correct. The externality at the subsidized equilibrium is smaller than at the unsubsidized one. By ignoring this feedback, the Pigouvian formula prescribes too much correction. The quantitative magnitude (23–32% overshoot, depending on group and steady state)

and its consequences for transition-path welfare are documented in Section 5.

4.3 Distortionary Financing

The lump-sum financing assumption is a useful benchmark, but governments in practice raise revenue through distortionary instruments. When each dollar of public spending costs society more than one dollar due to deadweight losses, the **marginal cost of public funds** (MCPF) is $1 + \mu$, where $\mu > 0$ captures the efficiency loss. This introduces a wedge between the externality benefit and the fiscal cost of the subsidy.

Modified planner's problem. The planner now maximizes welfare net of the social cost of financing. Here W is evaluated with $T = 0$ in agents' budget constraints, and the MCPF penalty $(1 + \mu)$ fully captures the social cost of raising revenue:

$$\max_{\tau_H, \tau_L} \; W(\tau_H, \tau_L) - (1 + \mu) \sum_i \lambda_i \tau_i n_i^*(\tau_i). \tag{26}$$

At $\boldsymbol{\tau} = \mathbf{0}$, a positive subsidy for group i improves welfare if

$$\frac{\lambda_i}{c_i} + \frac{\text{MEB}_i}{\phi_i} \;>\; (1 + \mu)\lambda_i. \tag{27}$$

The left-hand side is the marginal social benefit of an additional child from group i: the private benefit λ_i / c_i (transfer value at marginal utility of income) plus the externality benefit MEB_i / ϕ_i (external welfare gain per dollar of subsidy). The right-hand side is the marginal fiscal cost $(1 + \mu)\lambda_i$. A subsidy improves welfare only if the benefit exceeds the cost.

Proposition 4.3 (MCPF Threshold)**.** *A positive subsidy for group i is welfare-improving if and only if*

$$\mu < \mu_i^* \equiv \frac{\text{MEB}_i}{\lambda_i \phi_i} + \frac{1}{c_i} - 1. \tag{28}$$

When $\mu \geq \mu_i^*$, the deadweight loss from financing exceeds the welfare gain, and no subsidy is warranted for group i.

The proof is a rearrangement of (27) and is given in Appendix A.

The MCPF decomposition. The threshold μ_i^* has two components with distinct economic content:

$$\mu_i^* = \underbrace{\frac{\text{MEB}_i}{\lambda_i \phi_i}}_{\text{externality component}} + \underbrace{\left(\frac{1}{c_i} - 1\right)}_{\text{generic transfer component}}. \tag{29}$$

The generic transfer component $1/c_i - 1$ is the welfare gain from any lump-sum transfer to agents whose marginal utility of income $1/c_i$ exceeds unity. This term would be positive even without an externality: transferring a dollar to agents who spend only $c_i < 1$ on consumption raises welfare. The externality component $\text{MEB}_i / (\lambda_i \phi_i)$ is specific to the pronatal subsidy: it captures the additional welfare gain from directing the transfer toward children, which generates the empathy externality.

The decomposition answers a natural question: why not just give cash? A generic cash transfer captures only the transfer component. The baby bonus captures both. The ratio of the externality component to the generic component, $\text{MEB}_i / [\lambda_i \phi_i (1/c_i - 1)]$, measures the premium that the pronatal instrument earns over unconditional transfers. At calibrated parameters, this premium is substantial. The externality increases the MCPF headroom by approximately 72% relative to cash for Group H, and by about 40% for Group L. The quantitative details are in Section 5.

Relationship to lump-sum optimum. The distortionary-financing optimum τ_i^{**} satisfies $\tau_i^{**} < \tau_i^*$ (the lump-sum Pigouvian level), since the MCPF makes each dollar of subsidy costlier. The relationship is monotone: as μ increases from zero, the optimal subsidy shrinks continuously until it reaches zero at $\mu = \mu_i^*$. The empirical public finance literature estimates MCPF in the range $\mu \in [0.1, 0.5]$ for developed economies, with estimates varying by tax instrument and country (see, e.g., Browning, 1976 for the U.S. income tax). Whether the externality survives this range is a quantitative question addressed in Section 5.

4.4 Targeting

Given a budget for pronatal subsidies, which group should receive the larger allocation? Conventional wisdom suggests targeting the group with the largest private fertility response, that is, the group where each dollar of subsidy produces the most additional births. We show that this criterion can be incorrect.

Marginal welfare per fiscal dollar. The welfare gain per dollar spent on subsidizing group i, evaluated at $\boldsymbol{\tau} = \mathbf{0}$ with distortionary financing, is

$$\mathrm{MW}_i = \frac{1}{(1+\mu)\, c_i} + \frac{\mathrm{MEB}_i}{(1+\mu)\, \lambda_i\, \phi_i}. \tag{30}$$

The first term is the private transfer value per fiscal dollar. The second is the externality value per fiscal dollar. When consumption levels are similar across groups ($c_H \approx c_L$, as at calibrated parameters), the private terms approximately cancel and the ranking is determined by the externality-per-dollar ratios.

Proposition 4.4 (Targeting by Externality Rank)**.** *When consumption levels are similar across groups ($c_i \approx c_j$), the planner should allocate the marginal dollar to group i over group j whenever*

$$\frac{\mathrm{MEB}_i}{\lambda_i \phi_i} > \frac{\mathrm{MEB}_j}{\lambda_j \phi_j}, \tag{31}$$

that is, whenever group i has a higher externality rank (Definition 3.5). The optimal target is the group whose externality per dollar of subsidy is largest, not the group whose externality per child is largest.

The proof is immediate from comparing MW_i and MW_j in (30) when $c_i \approx c_j$. The full statement, accommodating consumption differences, is in Appendix A.

Intuition: the Ramsey analogy. The targeting principle has a direct analog in optimal taxation. The Ramsey rule says to tax goods with inelastic demand more heavily, because the deadweight loss per dollar of revenue is smallest. Our rule says to subsidize the group where the externality per fiscal dollar is highest. In both cases, the optimal

allocation depends on the ratio of the relevant margin to the fiscal cost, not on the absolute size of the margin alone.

Why the ranking can reverse. If Group L generates a larger MEB per child (because its larger population share gives it more exposure propagation) but also has a larger $\lambda_i \phi_i$ (because it is numerous and its children are expensive), the externality-per-dollar ranking can favor Group H. This is what occurs at calibrated parameters. A government that directs pronatal transfers toward the group with the most "visible" fertility shortfall may be misallocating resources. The correct criterion is the externality per fiscal dollar, which favors groups that are small, cohesive, and low-cost to subsidize.

4.5 Uniform vs. Group-Specific

Political or administrative constraints may require a uniform subsidy ($\tau_H = \tau_L = \tau$). The planner then solves a constrained version of (26):

$$\max_{\tau} \; W(\tau) - (1+\mu)\,\tau \sum_i \lambda_i \, n_i^*(\tau). \tag{32}$$

The optimal uniform rate is a weighted average of the group-specific optima, with weights reflecting population shares and fertility responses. When the externality-per-dollar ratios differ across groups, imposing uniformity entails a welfare loss.

Proposition 4.5 (Welfare Loss from Uniformity)**.** *When the externality-per-dollar ratios differ across groups, the uniform subsidy is strictly dominated by the group-specific subsidy.*

The formal proof follows the standard Ramsey logic. At the group-specific optimum (τ_H^*, τ_L^*), each group's first-order condition holds separately. The uniform constraint $\tau_H = \tau_L$ forces both groups onto the same rate, generating a second-order welfare loss proportional to $(\tau_H^* - \tau_L^*)^2$, which is positive whenever the externality-per-dollar ratios differ across groups. See Appendix A for the derivation.

Practical relevance. The welfare loss from uniformity is second-order in the gap between group-specific optima. At calibrated parameters (Section 5), the uniform subsidy captures over 98% of the group-specific welfare gain, so the practical recommendation leans toward uniformity. The group-specific optimum concentrates subsidies on Group H, generating a distributional cost that the uniform subsidy ameliorates. Section 5 examines these distributional effects, including the feasibility of Pareto-improving designs.

Four insights emerge from the welfare analysis. A positive externality justifies a subsidy, the Pigouvian formula provides a first approximation, the MCPF threshold disciplines the fiscal case, and the targeting rule allocates across groups.

5 Quantitative Analysis

This section calibrates the model to match pre-shock fertility differentials and quantifies the empathy channel's contribution to the fertility decline, the optimal subsidy, and its robustness. The empathy weight ψ governs the strength of the channel but cannot be pinned by a single calibration target. We set $\psi = 0.20$ as a conservative baseline and present results across $\psi \in [0.05, 0.50]$.

5.1 Calibration

Table 1 summarizes the model parameters, organized by source.

Calibration strategy. Panel A parameters are derived from de la Croix and Doepke (2003): $\theta = 0.372$ matches their child-to-consumption utility ratio (0.271/0.729), and per-child costs are set so the population-weighted average ($\bar{\phi} = 0.075$) matches their benchmark, with $\phi_H/\phi_L = 0.75$ reflecting lower opportunity costs in communal child-rearing environments. The population share $\lambda_H = 0.25$ falls within the range of diary-corrected weekly attendance estimates (20–30%; Brenner, 2011). Panel B parameters are sensitivity-tested in Section 5.4. Results are insensitive to δ and ϕ_H/ϕ_L but sensitive to

Table 1. *Model Parameters*

Parameter	Symbol	Value	Source
Panel A: Literature			
Preference weight	θ	0.372	de la Croix and Doepke (2003)
Per-child cost (H)	ϕ_H	0.060	Wtd. avg. matches de la Croix and Doepke (2003)
Per-child cost (L)	ϕ_L	0.080	$\phi_H/\phi_L = 0.75$
Population share (H)	λ_H	0.25	Diary-corrected attendance range 20–30% (Brenner, 2011)
Panel B: Assumption (tested)			
Segregation (H/L)	δ_H/δ_L	0.70/0.30	Community structure
Empathy rate	ρ	1.0	$e_H \approx 0.94$ at pre-shock steady state
Quality cost	γ	0.06	$h_L \approx 0.76$ at baseline
Discount factor	β	0.36	0.96^{25} (planner's present value only)
Panel C: Novel			
Empathy weight	ψ	0.20	Swept $[0.05, 0.50]$
Panel D: Calibrated			
Quantity weight (H/L)	ν_H/ν_L	0.587/0.603	Match $n_H = 3.0$, $n_L = 2.3$
Panel E: Shock			
Cost shock (H/L)	Δ_H/Δ_L	17%/23%	Skill-biased tech. change; $\Delta_L/\Delta_H \approx 1.35$

Notes. Only ν_H and ν_L are calibrated endogenously. The empathy weight ψ is set at a conservative baseline and swept over $[0.05, 0.50]$ in Section 5.4. Income normalized to $w = 1$.

ρ and ψ. The empathy weight $\psi = 0.20$ (Panel C) is set conservatively: at this value, the model's social multiplier is 1.10–1.12×, well below the large dyadic effects estimated in the peer-effects literature (Lyngstad and Prskawetz, 2010; Balbo and Barban, 2014), which capture all social interaction channels (conformity, learning, and empathy jointly) while ψ governs only empathy.

Calibrated and shock parameters. The quantity weights $\nu_H = 0.587$ and $\nu_L = 0.603$ (Panel D) are the only endogenous parameters, pinned by matching pre-shock fertility targets: $n_H = 3.0$ and $n_L = 2.3$. These are illustrative targets chosen to represent a stylized pre-shock steady state with a meaningful fertility differential between groups.

They are informed by General Social Survey (GSS) microdata on completed fertility by religious attendance (GSS 1972–2024, authors' calculations), where children ever born (CEB) for weekly-or-more attenders ranges from 2.5 to 3.2 and for the rest of the population from 1.9 to 2.7, depending on birth cohort and survey era. The qualitative results of the model hold across a wide range of targets, and Section 5.4 confirms robustness to parameter variation. Since ν_i absorbs the effect of ψ, the empathy share of the fertility decline is determined by the choice of ψ, not estimated from the data. The cost shock (Panel E) is motivated by skill-biased technological change: child-rearing costs rise by 17% for Group H and 23% for Group L ($\phi_H^{\text{post}} = 0.0702$, $\phi_L^{\text{post}} = 0.0984$), with the larger increase for Group L reflecting the greater opportunity-cost increase for less cohesive communities ($\Delta_L / \Delta_H \approx 1.35$).

Baseline equilibrium. At the pre-shock steady state, both groups maintain positive quality investment ($h_H = 0.606$, $h_L = 0.761$), confirming that the model generates an interior quantity-quality tradeoff. Consumption is equal across groups ($c_H = c_L = 0.711$). The empathy signal is near saturation ($e_H = 0.942$, $e_L = 0.911$), so the channel is largely dormant at pre-shock fertility, and intrinsic preferences account for more than half of the pro-quantity motive in both groups.

5.2 Cost Shock and the Empathy Decomposition

The skill-biased technological change (SBTC) shock raises child-rearing costs and triggers a fertility decline whose anatomy is shown in Table 2.

The decline and decomposition. Aggregate fertility falls from $\bar{N} = 2.475$ to $\bar{N} = 1.977$ (20.1%), with the decline larger for Group L (21.7%) than Group H (16.6%), reflecting the asymmetric cost shock. Both groups increase quality investment (h_H: $0.606 \to 0.731$; h_L: $0.761 \to 0.980$), confirming that the quantity-quality substitution is operative. To isolate the empathy channel, we compare the full model to a $\psi = 0$ counterfactual (recalibrating ν_i to match the same pre-shock targets). The counterfactual decline is

Table 2. *Equilibrium, Cost Shock, and Empathy Decomposition*

	n_H	n_L	h_H	h_L	$\bar{c}$	$\bar{N}$
Panel A: Pre-shock steady state						
Full model ($\psi = 0.20$)	3.000	2.300	0.606	0.761	0.711	2.475
Panel B: Post-shock steady state						
Full model ($\psi = 0.20$)	2.501	1.802	0.731	0.980	0.716	1.977
Counterfactual ($\psi = 0$)	2.564	1.870	—	—	—	2.043
Panel C: Decomposition of $\bar{N}$ decline						
Total decline (full model)					0.498	(20.1%)
Cost-only decline ($\psi = 0$)					0.432	(17.4%)
Empathy amplification					0.067	(13.4%)
Empathy share range ($\psi \in [0.05, 0.50]$)						[3%, 33%]

Notes. The counterfactual $\psi = 0$ model recalibrates ν_i to match the same pre-shock fertility targets, so the decomposition isolates the role of the empathy feedback. $\bar{c}$ denotes the population-weighted average consumption. In Panel C, the total and cost-only decline percentages are shares of the pre-shock $\bar{N}$. The empathy amplification percentage (13.4%) is the empathy share of the total decline (0.067/0.498). All Panel C values are computed at full precision before rounding.

0.432 (17.4%), and the empathy amplification is 0.067 children per woman, or 13.4% of the total. Across $\psi \in [0.05, 0.50]$, the empathy share spans 3% to 33%. This range is the appropriate summary. The 13.4% figure at $\psi = 0.20$ illustrates the mechanism at a plausible value. The entire amplification operates through the feedback loop ($n \to E \to e \to n$): fixing the empathy signal at its pre-shock level produces the same long-run fertility as setting $\psi = 0$.

Figure 2 shows the adjustment path. Convergence is fast (2–3 generations, reflecting the spectral radius of 0.123), and the empathy gap widens over time as feedback accumulates. At the pre-shock steady state, both groups sit in the near-saturated region ($e > 0.91$), where marginal sensitivity is small, but as fertility falls, the operating point moves into the steep region and the channel "wakes up." This explains both why the amplification is modest at baseline and why it grows with larger cost shocks.

5.3 GE-Optimal Subsidies

This subsection computes GE-optimal subsidies by numerical welfare maximization and compares them to the Pigouvian benchmark.

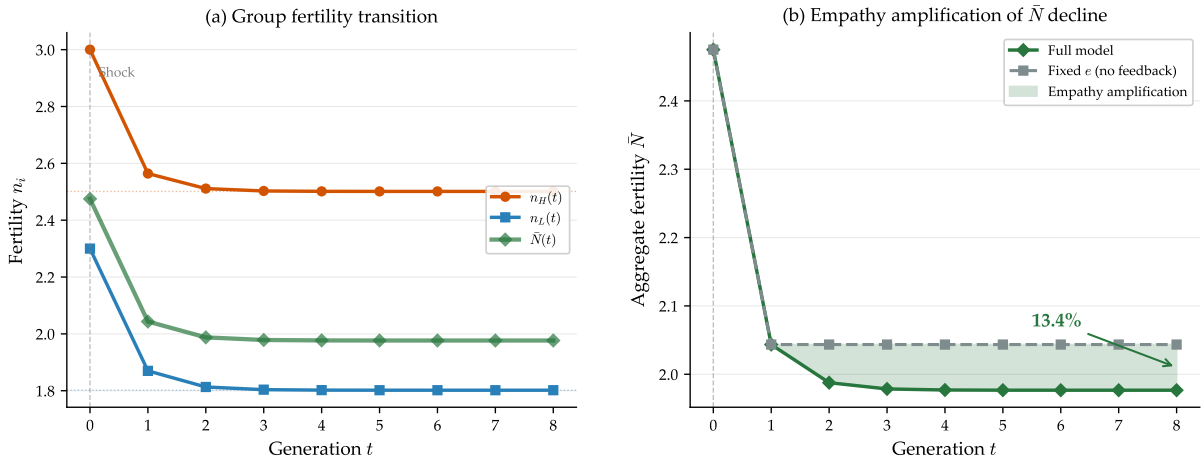

Figure 2. *Transition dynamics and empathy decomposition. Panel (a): group-specific fertility $n_H(t)$ and $n_L(t)$ after the cost shock, converging to the new steady state in 2–3 generations. Panel (b): aggregate fertility $\bar{N}$ in the full model (solid) versus the $\psi = 0$ counterfactual (dashed). The shaded gap is the empathy amplification, which accumulates to 13.4% of the total decline at the new steady state.*

Pigouvian overshoot. Table 3 compares four subsidy schemes against the no-subsidy baseline at the pre-shock steady state. The Pigouvian rates ($\tau_H = 0.01257$, $\tau_L = 0.00935$) exceed the GE-optimal rates ($\tau_H^* = 0.00952$, $\tau_L^* = 0.00760$) by 32% and 23%, respectively. A planner who implements the Pigouvian formula overspends: the fiscal cost is 30% higher, the induced fertility overshoot pushes the empathy signal into the saturated region, and the transition-path present value turns negative.

The welfare surface. The welfare function $W(\tau_H, \tau_L)$ is relatively flat near the optimum (Figure 3). The GE-optimal point sits closer to the origin than the Pigouvian point, and the modest gap between uniform and group-specific GE-optimal rates translates into a welfare difference of $+0.000031$.

Group-specific dominates uniform (but barely). At the GE optimum, group-specific targeting yields $\Delta W = +0.00219$ versus $+0.00216$ for uniform: the advantage is $+0.000031$, or 1.4% of the total gain. The uniform subsidy captures over 98% of the welfare improvement. The earlier finding that the uniform Pigouvian appeared to outperform the group-specific Pigouvian was an artifact of comparing both at their respective overshooting rates. At GE-optimal rates, the expected ranking is restored. A

Table 3. *Policy Dashboard: Subsidy Schemes at the Pre-Shock Steady State*

Scheme	τ_H	τ_L	$\tilde{N}$	ΔW	CE (%)
No subsidy	0	0	2.475	—	—
First-order Pigouvian					
Group-specific	0.01257	0.00935	2.863	+0.00197	0.20
Uniform	0.01015	0.01015	2.840	+0.00206	0.21
GE-optimal					
Uniform	0.00845	0.00845	2.773	+0.00216	0.22
Group-specific	0.00952	0.00760	2.771	**+0.00219**	**0.22**

Notes. ΔW is the change in steady-state utilitarian welfare relative to no subsidy. CE is the consumption-equivalent welfare gain. The GE-optimal rates are obtained by numerically maximizing $W(\tau_H, \tau_L)$ over the full subsidized equilibrium. The Pigouvian group-specific rates overshoot the GE-optimal by 32% (τ_H) and 23% (τ_L). The GE-optimal targeting ratio is $\tau_H^*/\tau_L^* = 1.25$.

single subsidy rate captures nearly all of the welfare gain while avoiding differentiated transfers.

Welfare magnitudes. The GE-optimal subsidy raises welfare by 0.22% CE. Using U.S. median household income ($75,000) and $\bar{c} \approx 0.71$, this translates to approximately $117 per family per year.[1] At $\psi = 0.40$, the welfare gain increases roughly 3.5-fold.

The MCPF thresholds ($\mu_H^* = 0.70$, $\mu_L^* = 0.57$) exceed typical estimates of deadweight loss ($\mu \in [0.1, 0.5]$), confirming that the subsidy survives distortionary financing. The threshold decomposes into a generic transfer component ($1/c_i - 1 = 0.407$) and an externality component ($\text{MEB}_i / (\lambda_i \phi_i) = 0.295$ for H, 0.164 for L). The externality allows a baby bonus to withstand 72% higher fiscal costs than an equivalent cash transfer before becoming welfare-negative for Group H, because it captures both the transfer value and the externality correction.

The targeting reversal. Figure 4 illustrates the key policy finding. Group L generates a larger externality per child ($\text{MEB}_L = 0.00986 > \text{MEB}_H = 0.00442$), but the optimal

[1] Illustrative: $0.0022 \times 0.71 \times \$75{,}000 \approx \$117$. The mapping from model to U.S. incomes is approximate.

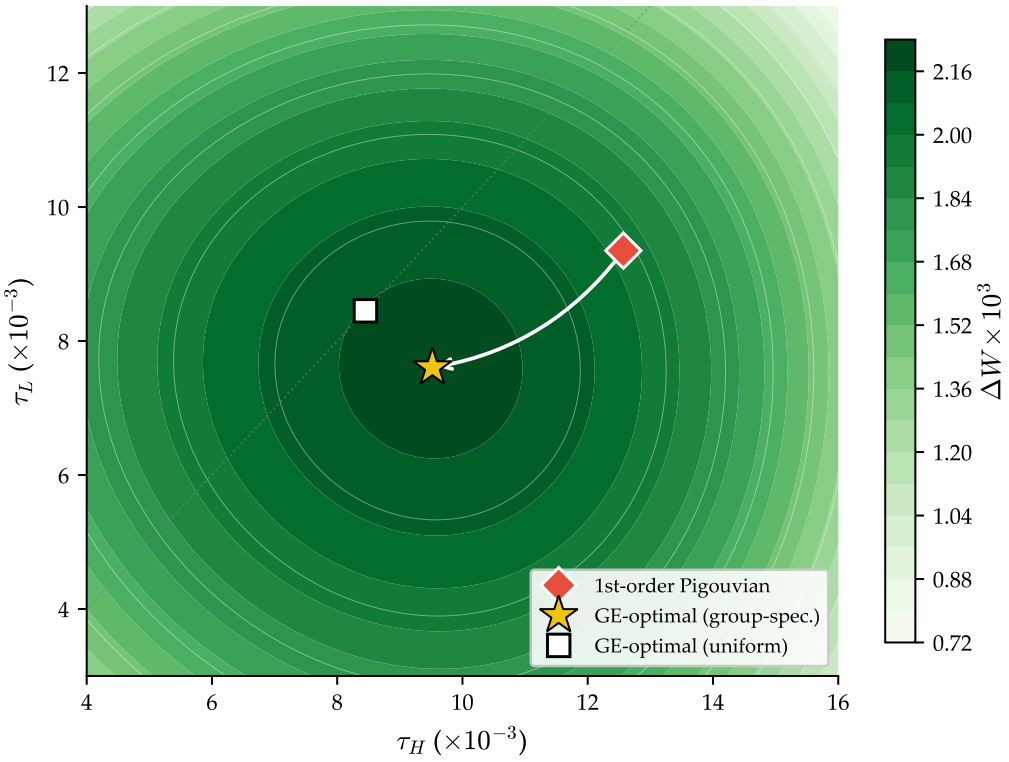

Figure 3. *Welfare surface in (τ_H, τ_L) space. Darker regions indicate higher welfare gain. The first-order Pigouvian rates (diamond) overshoot the GE-optimal (star) by 23–32%, landing in a lower-welfare region. The GE-optimal uniform rate (square) lies on the 45-degree line. The welfare function is approximately flat near the optimum, explaining the small gap between uniform and group-specific GE-optimal schemes.*

target is H because the relevant criterion is the externality per fiscal dollar:

$$\frac{\text{MEB}_H}{\lambda_H \phi_H} = \frac{0.00442}{0.25 \times 0.06} = 0.295 \quad > \quad \frac{\text{MEB}_L}{\lambda_L \phi_L} = \frac{0.00986}{0.75 \times 0.08} = 0.164.$$

Group H is a smaller population with lower per-child costs, so each fiscal dollar produces more empathy spillover. The GE-optimal subsidy ratio $\tau_H^*/\tau_L^* = 1.25$ reflects this.

The post-shock case. The case for intervention strengthens after the cost shock. At the post-shock steady state ($\bar{N} = 1.977$), the MEB rises by 28% for Group H and 18% for Group L, the GE-optimal welfare gain increases from $+0.00219$ to $+0.00265$, and the targeting ratio rises from 1.25 to 1.36. As fertility falls, the externality grows and the scope for welfare-improving intervention expands.

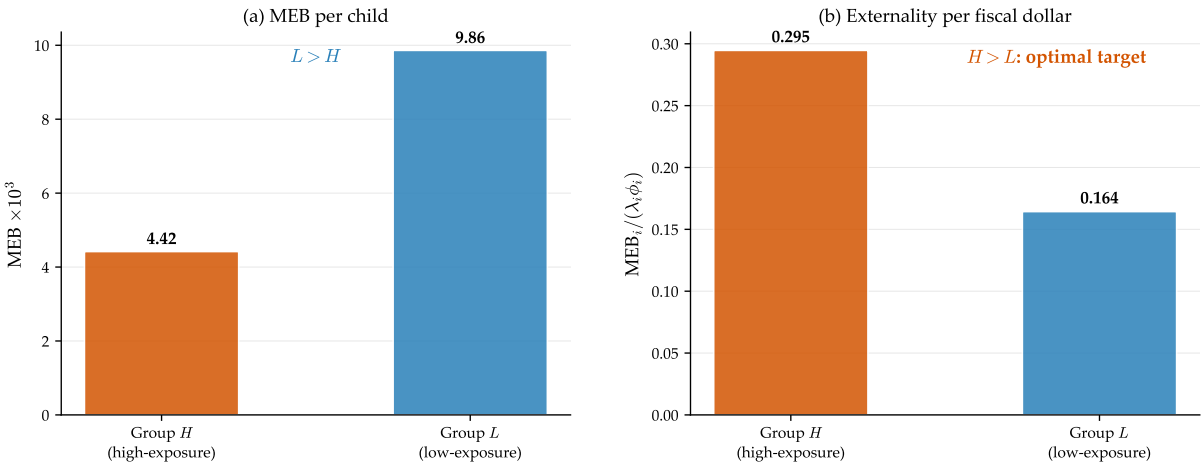

Figure 4. *The targeting reversal. Panel (a): Group L generates a larger marginal external benefit per child* ($\text{MEB}_L > \text{MEB}_H$)*. Panel (b): after dividing by population share and per-child cost, Group H generates a larger externality per dollar of subsidy* ($\text{MEB}_H / (\lambda_H \phi_H) > \text{MEB}_L / (\lambda_L \phi_L)$)*. The optimal target is H, the group where each fiscal dollar produces the most externality.*

Transition-path welfare. At GE-optimal rates with $\beta = 0.36$ ($= 0.96^{25}$), the transition present value (PV) is approximately zero ($+0.00002$): the front-loaded fiscal cost and the gradually accumulating empathy benefits nearly offset. The subsidy is a long-run instrument. At a lower discount rate ($\beta = 0.50$), the PV turns clearly positive. The Pigouvian rates produce a negative PV (-0.00073) because overshooting imposes a larger fiscal burden than the welfare gain justifies.

Distributional effects. Group L bears a net fiscal cost under all subsidy schemes. Under group-specific targeting, it receives a smaller per-child subsidy than Group H but pays the same lump-sum tax. The externality benefit Group L receives from Group H's additional children does not fully offset this fiscal drain. Under the uniform scheme, Group L's loss is 6.5 times smaller than under group-specific targeting (-0.0004 versus -0.003). A Pareto-improving design exists (mildly inverted targeting: $\tau_L = 0.00875 > \tau_H = 0.00825$), at a cost of only 2.2% of the total welfare improvement.

5.4 Robustness

The key questions for robustness are whether the qualitative results (positive externality, subsidy case, targeting rule) survive across the parameter space, and how the

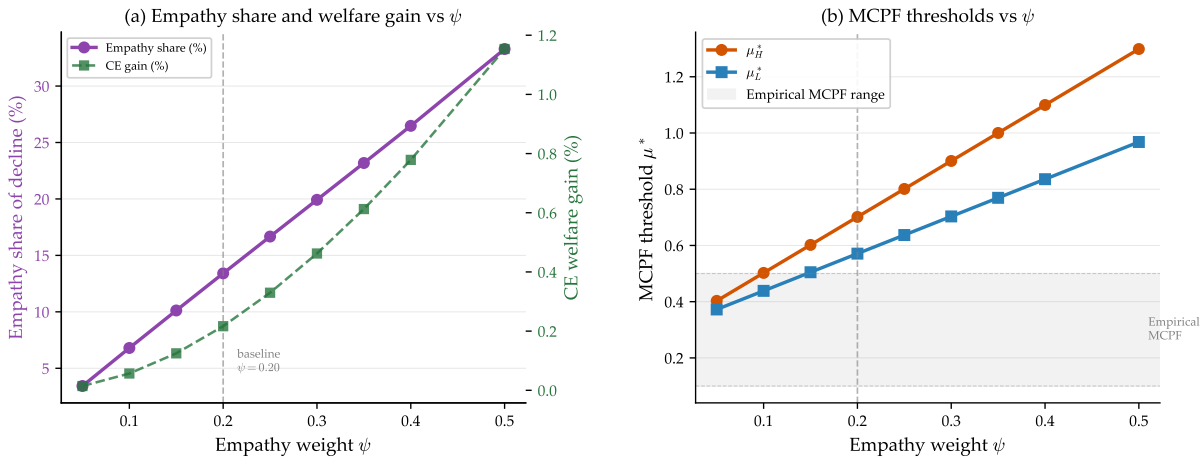

Figure 5. *Robustness to the empathy weight ψ. Panel (a): the empathy share of the fertility decline (left axis) scales monotonically with ψ. The GE-optimal CE welfare gain (right axis) rises monotonically with ψ. The baseline $\psi = 0.20$ is marked. Panel (b): the MCPF threshold μ^* for both groups. The shaded band shows the empirical MCPF range ($\mu \in [0.1, 0.5]$). Both groups' thresholds exceed the upper end of this range for $\psi \geq 0.20$.*

magnitudes depend on the choice of ψ. Figure 5 presents the central robustness exercise, sweeping ψ over $[0.05, 0.50]$ with ν_i recalibrated at each point to match the same pre-shock fertility targets.

Sensitivity to ψ. The empathy share, MCPF thresholds, and social multiplier all scale monotonically with the empathy weight. The empathy share ranges from 3.4% ($\psi = 0.05$) to 33.3% ($\psi = 0.50$), and the MCPF threshold for Group H rises from 0.40 to 1.30. At $\psi \geq 0.20$, both groups' thresholds exceed the upper end of the empirical MCPF range ($\mu = 0.5$), while at $\psi = 0.05$–0.10, the subsidy may not survive at the highest MCPF estimates. The social multiplier spans 1.02–1.32$\times$ across the sweep, a range conservative relative to the peer-effects literature (Lyngstad and Prskawetz, 2010; Balbo and Barban, 2014).

Structural parameters. The segregation parameters and cost ratio have negligible effects: sweeping δ_H across $[0.3, 0.9]$ changes the empathy share by about one percentage point. The empathy rate ρ has a quantitatively significant effect: at $\rho = 0.50$, the empathy share reaches 21.6%, while at $\rho = 2.0$ it falls to 2.7%. The sensitivity arises because ψ and ρ jointly govern the marginal externality at the pre-shock exposure level.

Full robustness tables are in Appendix B.

The qualitative results are robust across the parameter space. The magnitudes scale with ψ and ρ in predictable ways. The empathy share of the fertility decline ranges from 3% to 33%, the optimal subsidy raises welfare by 0.22% CE at the baseline $\psi = 0.20$, and the subsidy survives distortionary financing for $\psi \geq 0.20$.

5.5 Transition-Optimal Policy

The GE-optimal subsidies computed above maximize steady-state welfare. During the transition, however, the economy has not yet reached its new steady state, and the planner may wish to set a different subsidy for each generation. A natural conjecture is that subsidies should be front-loaded, since the economy is furthest from its new steady state early in the transition. We investigate this conjecture by allowing the planner to choose a time-varying path $\{\tau_i(t)\}_{t=0}^{T}$ and find, perhaps surprisingly, that the optimal path is approximately flat and substantially below the post-shock steady-state optimum.

Consider a planner who chooses per-child subsidies $\tau_H(t)$ and $\tau_L(t)$ for each generation $t = 0, 1, \ldots, T$ to maximize the present value of social welfare along the transition, subject to a balanced budget in each period. The economy starts at the pre-shock steady state $(n_H, n_L) = (3.0, 2.3)$ and faces post-shock costs $(\phi_H^{\text{post}}, \phi_L^{\text{post}})$ from $t = 0$ onward. Each generation's choices depend on the predetermined child stock (shaped by the previous generation's subsidized fertility), with lump-sum balanced-budget financing. The planner discounts future welfare at the generational rate $\beta = 0.36$.

Applying the post-shock steady-state optimal rates during the transition produces a *negative* present value relative to no subsidy ($\Delta\text{PV} \approx -0.0009$ for both uniform and group-specific schemes). The mechanism is fiscal. At the start of the transition, fertility is still near its pre-shock level ($\bar{N} = 2.475$), so the per-period fiscal cost is high. With $\beta = 0.36$, these early fiscal burdens receive large weight in the present value, while the welfare gains from the externality correction accumulate slowly and are heavily discounted by the time the economy reaches the new steady state. Rates calibrated for the long run are too generous for the transition.

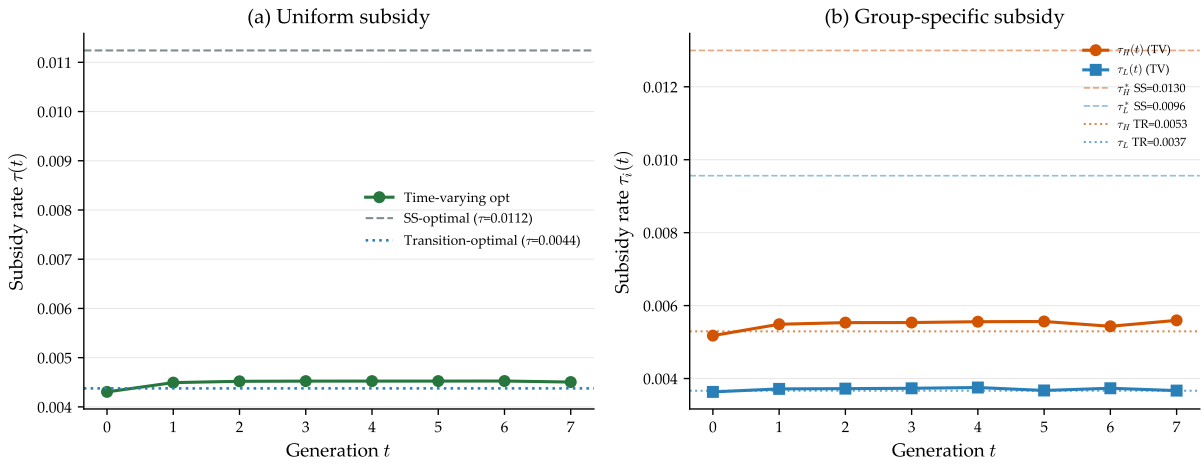

Figure 6. *Optimal subsidy paths during the post-shock transition. Panel (a): uniform subsidy. Panel (b): group-specific subsidies. In each panel, the dashed line shows the steady-state optimal rate, the dotted line shows the transition-optimal constant rate (≈40% of the steady-state rate), and the solid line shows the time-varying optimal path. The time-varying path is approximately flat, confirming that the welfare gain comes from level adjustment, not from temporal variation.*

Optimizing over constant-rate subsidies applied throughout the transition yields rates approximately 40% of the steady-state optimum ($\tau \approx 0.004$ uniform, $\tau_H \approx 0.005$ and $\tau_L \approx 0.004$ group-specific). To assess whether temporal variation can improve on this, we decompose the welfare gain from the full time-varying solution into two components: the gain from adjusting the *level* of the subsidy (holding the path flat at the transition-optimal constant) and the gain from allowing the *shape* to vary across generations. Essentially all of the improvement comes from the level adjustment, and the contribution of temporal variation is negligible (Figure 6). Group-specific targeting retains its value along the transition, with the time-varying path maintaining $\tau_H(t)/\tau_L(t) \approx 1.4$, comparable to the steady-state targeting ratio.

Why is the optimal path flat rather than front-loaded? Two forces push against early generosity. First, at the start of the transition, the empathy signal remains near saturation ($e_H = 0.94$, $e_L = 0.91$), so the marginal externality that the subsidy corrects is small. As fertility falls and exposure declines into the steeper region of the empathy function, the marginal return to correction rises, favoring later rather than earlier intervention. Second, the budget cost per dollar of subsidy is proportional to fertility, which starts high. Early subsidies are expensive precisely because there are many

children to cover. At the baseline $\beta = 0.36$, the resulting front-loading ratio (the ratio of the first-period rate to the last) is approximately 0.96, indicating a nearly uniform path. Sensitivity analysis confirms that front-loading only emerges for $\beta \geq 0.70$, where the planner values future welfare gains enough to justify the early fiscal outlay, and even then the welfare gain from temporal variation remains negligible.

These results suggest that the per-child subsidy is a long-run instrument. During demographic transitions, applying roughly 40% of the steady-state rate as a flat subsidy captures the available welfare gain, with no meaningful improvement from temporal variation. Targeting, by contrast, remains valuable throughout: directing a larger subsidy toward the high-exposure group raises the present value of welfare along the transition, just as it does at the steady state.

6 Conclusion

This paper has embedded a biologically grounded empathy channel in a two-group quantity-quality model of fertility. Exposure to infants in the social environment activates an affective response that increases the desire for parenthood, and this response generates a positive externality that the decentralized equilibrium fails to internalize. The concavity of the empathy function implies that the channel is largely dormant at high fertility but activates as fertility falls, amplifying cost-driven declines through a feedback loop.

On the policy side, the externality justifies a per-child subsidy on standard economic grounds. However, the first-order Pigouvian correction overshoots the general-equilibrium optimum because the subsidy itself raises fertility and partially closes the externality gap. The optimal targeting rule follows Ramsey logic, directing resources toward the group where each fiscal dollar generates the most externality rather than the group with the largest per-child externality. The subsidy survives distortionary financing at realistic estimates of the marginal cost of public funds. A uniform subsidy captures nearly all of the welfare gain from group-specific targeting, offering a practical

design that avoids differentiated transfers. During demographic transitions, the optimal rate is well below the steady-state optimum, because the fiscal cost of the subsidy is proportional to fertility, which starts high.

The welfare case rests on including empathy-modified preferences in the social welfare function. We have argued that this is warranted on evolutionary, behavioral, and welfarist grounds, and that the policy corrects an externality rather than manipulating preferences. Alternative normative positions are possible, and the welfare conclusions are sensitive to this choice.

More broadly, the model suggests that fertility motivation is not a fixed parameter but responds to the demographic environment. As fertility declines and infant exposure falls, the conditions that sustain the desire for children weaken, and the decline can become self-reinforcing. Modern patterns of residential segregation by age, privatized childcare, and reduced community contact with infants may have eroded a source of fertility motivation that was once widespread. Whether this channel is quantitatively important is ultimately an empirical question, but the framework developed here provides the tools to ask it.

References

Abraham, Eyal, Talma Hendler, Irit Shapira-Lichter, Yaniv Kanat-Maymon, Orna Zagoory-Sharon, and Ruth Feldman, "Father's Brain Is Sensitive to Childcare Experiences," *Proceedings of the National Academy of Sciences*, 2014, *111* (27), 9792–9797.

Balbo, Nicoletta and Nicola Barban, "Does Fertility Behavior Spread among Friends?," *American Sociological Review*, 2014, *79* (3), 412–431.

Becker, Gary S., "An Economic Analysis of Fertility," in "Demographic and Economic Change in Developed Countries," Columbia University Press, 1960, pp. 209–240. NBER Conference Series 11.

_ **and H. Gregg Lewis**, "On the Interaction between the Quantity and Quality of Children," *Journal of Political Economy*, 1973, *81* (2), S279–S288.

Bernheim, B. Douglas and Antonio Rangel, "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics," *Quarterly Journal of Economics*, 2009, *124* (1), 51–104.

Bisin, Alberto and Thierry Verdier, "Beyond the Melting Pot: Cultural Transmission, Marriage, and the Evolution of Ethnic and Religious Traits," *Quarterly Journal of Economics*, 2000, *115* (3), 955–988.

_ **and** _ , "The Economics of Cultural Transmission and the Dynamics of Preferences," *Journal of Economic Theory*, 2001, *97* (2), 298–319.

Brase, Gary L. and Sandra L. Brase, "Emotional Regulation of Fertility Decision Making: What Is the Nature and Structure of "Baby Fever"?," *Emotion*, 2012, *12* (5), 1141–1154.

Brenner, Philip S., "Identity Importance and the Overreporting of Religious Service Attendance: Multiple Imputation of Religious Attendance Using the American Time Use Survey and the General Social Survey," *Journal for the Scientific Study of Religion*, 2011, *50* (1), 103–115.

Brock, William A. and Steven N. Durlauf, "Discrete Choice with Social Interactions," *Review of Economic Studies*, 2001, *68* (2), 235–260.

_ **and** _ , "Multinomial Choice with Social Interactions," in Lawrence E. Blume and Steven N. Durlauf, eds., *The Economy as an Evolving Complex System III*, Oxford University Press, 2006, pp. 175–206.

Browning, Edgar K., "The Marginal Cost of Public Funds," *Journal of Political Economy*, 1976, *84* (2), 283–298.

Burkart, Judith M., Sarah B. Hrdy, and Carel P. van Schaik, "Cooperative Breeding and Human Cognitive Evolution," *Evolutionary Anthropology*, 2009, *18* (5), 175–186.

Büyükkeçeci, Zafer, Thomas Leopold, Ruben van Gaalen, and Henriette Engelhardt, "Family, Firms, and Fertility: A Study of Social Interaction Effects," *Demography*, 2020, *57* (1), 243–266.

de la Croix, David and Matthias Doepke, "Inequality and Growth: Why Differential Fertility Matters," *American Economic Review*, 2003, *93* (4), 1091–1113.

Doepke, Matthias and Fabian Kindermann, "Bargaining over Babies: Theory, Evidence, and Policy Implications," *American Economic Review*, 2019, *109* (9), 3264–3306.

_ **, Anne Hannusch, Fabian Kindermann, and Michèle Tertilt**, "The Economics of Fertility: A New Era," in Shelly Lundberg and Alessandra Voena, eds., *Handbook of the Economics of the Family*, Vol. 1, North-Holland, 2023, pp. 151–254.

Ericksen, Julia A., Eugene P. Ericksen, John A. Hostetler, and Gertrude E. Huntington, "Fertility Patterns and Trends among the Old Order Amish," *Population Studies*, 1979, *33* (2), 255–276.

Feldman, Ruth, "Oxytocin and Social Affiliation in Humans," *Hormones and Behavior*, 2012, *61* (3), 380–391.

_ , "The Neurobiology of Human Attachments," *Trends in Cognitive Sciences*, 2017, *21* (2), 80–99.

Fernandez, Raquel and Alessandra Fogli, "Fertility: The Role of Culture and Family Experience," *Journal of the European Economic Association*, 2006, *4* (2-3), 552–561.

_ **and** _ , "Culture: An Empirical Investigation of Beliefs, Work, and Fertility," *American Economic Journal: Macroeconomics*, 2009, *1* (1), 146–177.

Frejka, Tomas and Charles F. Westoff, "Religion, Religiousness and Fertility in the US and in Europe," *European Journal of Population*, 2008, *24* (1), 5–31.

Galiani, Sebastian and Raul A. Sosa, "Composition Beats Collapse: Insights from the Bisin–Verdier Model on Endogenous Fertility Reversal," Working Paper 34157, National Bureau of Economic Research 2025.

Galor, Oded and David N. Weil, "Population, Technology, and Growth: From Malthusian Stagnation to the Demographic Transition and Beyond," *American Economic Review*, 2000, *90* (4), 806–828.

Gauthier, Anne H., "The Impact of Family Policies on Fertility in Industrialized Countries: A Review of the Literature," *Population Research and Policy Review*, 2007, *26* (3), 323–346.

Glaeser, Edward L., Bruce Sacerdote, and Jose A. Scheinkman, "The Social Multiplier," *Journal of the European Economic Association*, 2003, *1* (2–3), 345–353.

Glocker, Melanie L., Daniel D. Langleben, Kosha Ruparel, James W. Loughead, Jeffrey N. Valdez, Mark D. Griffin, Norbert Sachser, and Ruben C. Gur, "Baby Schema Modulates the Brain Reward System in Nulliparous Women," *Proceedings of the National Academy of Sciences*, 2009, *106* (22), 9115–9119.

Goldin, Claudia, "Babies and the Macroeconomy," *Economica*, 2025, *92* (367), 675–700.

_ , "The Downside of Fertility," Working Paper 34268, NBER 2025.

_ **and Lawrence F. Katz**, "The Power of the Pill: Oral Contraceptives and Women's Career and Marriage Decisions," *Journal of Political Economy*, 2002, *110* (4), 730–770.

Greksa, Lawrence P., "Population Growth and Fertility Patterns in an Old Order Amish Settlement," *Annals of Human Biology*, 2002, *29* (2), 192–201.

Hickok, Gregory, "Eight Problems for the Mirror Neuron Theory of Action Understanding in Monkeys and Humans," *Journal of Cognitive Neuroscience*, 2009, *21* (7), 1229–1243.

Hrdy, Sarah Blaffer, *Mother Nature: A History of Mothers, Infants, and Natural Selection*, New York: Pantheon Books, 1999.

_ , *Mothers and Others: The Evolutionary Origins of Mutual Understanding*, Cambridge, MA: Harvard University Press, 2009.

Jones, Charles I., "The End of Economic Growth? Unintended Consequences of a Declining Population," *American Economic Review*, 2022, *112* (11), 3489–3527.

Kearney, Melissa S. and Phillip B. Levine, "Why Is Fertility So Low in High Income Countries?," Working Paper 33989, National Bureau of Economic Research 2025.

Kohler, Hans-Peter, Francesco C. Billari, and José Antonio Ortega, "The Emergence of Lowest-Low Fertility in Europe During the 1990s," *Population and Development Review*, 2002, *28* (4), 641–680.

Kramer, Karen L., "Cooperative Breeding and its Significance to the Demographic Success of Humans," *Annual Review of Anthropology*, 2010, *39*, 417–436.

Kringelbach, Morten L., Eloise A. Stark, Catherine Alexander, Marc H. Bornstein, and Alan Stein, "On Cuteness: Unlocking the Parental Brain and Beyond," *Trends in Cognitive Sciences*, 2016, *20* (7), 545–558.

Kuziemko, Ilyana, "Is Having Babies Contagious? Estimating Fertility Peer Effects between Siblings," Working Paper, Columbia University 2006.

Lesthaeghe, Ron, "The Unfolding Story of the Second Demographic Transition," *Population and Development Review*, 2010, *36* (2), 211–251.

Lorenz, Konrad, "Die angeborenen Formen möglicher Erfahrung," *Zeitschrift für Tierpsychologie*, 1943, *5* (2), 235–409.

Lutz, Wolfgang, Vegard Skirbekk, and Maria Rita Testa, "The Low-Fertility Trap Hypothesis: Forces that May Lead to Further Postponement and Fewer Births in Europe," *Vienna Yearbook of Population Research*, 2006, *4*, 167–192.

Lyngstad, Torkild Hovde and Alexia Prskawetz, "Do Siblings' Fertility Decisions Influence Each Other?," *Demography*, 2010, *47* (4), 923–934.

Manski, Charles F., "Identification of Endogenous Social Effects: The Reflection Problem," *Review of Economic Studies*, 1993, *60* (3), 531–542.

_ **and Joram Mayshar**, "Private Incentives and Social Interactions: Fertility Puzzles in Israel," *Journal of the European Economic Association*, 2003, *1* (1), 181–211.

Rotkirch, Anna, "All that She Wants Is A(nother) Baby? Longing for Children as a Fertility Incentive of Growing Importance," *Journal of Evolutionary Psychology*, 2007, *5* (1), 89–104.

_ **, Stuart Basten, Heini Väisänen, and Markus Jokela**, "Baby Longing and Men's Reproductive Motivation," *Vienna Yearbook of Population Research*, 2011, *9*, 283–306.

Sear, Rebecca and Ruth Mace, "Who Keeps Children Alive? A Review of the Effects of Kin on Child Survival," *Evolution and Human Behavior*, 2008, *29* (1), 1–18.

Singer, Tania and Claus Lamm, "The Social Neuroscience of Empathy," *Annals of the New York Academy of Sciences*, 2009, *1156*, 81–96.

Spolaore, Enrico and Romain Wacziarg, "Fertility and Modernity," *The Economic Journal*, 2022, *132* (642), 796–833.

Stigler, George J. and Gary S. Becker, "De Gustibus Non Est Disputandum," *American Economic Review*, 1977, *67* (2), 76–90.

Wasao, Samson, Cory Anderson, and Christian Mpody, "The Persistently High Fertility of a North American Population: A 25-Year Restudy of Parity among the Ohio Amish," *Population Studies*, 2021, *75* (3), 477–486.

A Proofs and Derivations

This appendix provides formal proofs for the propositions stated in the main text. Cross-references to equations use the main-text numbering.

A.1 Derivation of the Closed-Form Solution

The agent's problem (Section 3.4) is:

$$\max_{n_i>0,\, h_i>0} \ln\!\big(w - \phi_i n_i - \gamma n_i h_i\big) + \theta\big[\nu_i \ln(n_i) + (1-\nu_i)\ln(h_i)\big] + \psi e_i \ln(n_i)$$

with e_i taken as given by the agent.

Step 1: FOC for quality. Differentiating with respect to h_i and using $c_i = w - \phi_i n_i - \gamma n_i h_i$:

$$\frac{\theta(1-\nu_i)}{h_i} - \frac{\gamma n_i}{c_i} = 0 \qquad \Longrightarrow \qquad h_i = \frac{\theta(1-\nu_i)\, c_i}{\gamma\, n_i}$$

Step 2: Quality-spending identity. Multiplying both sides by γn_i:

$$\gamma n_i h_i = \theta(1-\nu_i)\, c_i$$

This is the Cobb-Douglas property: total quality spending is a fixed fraction of consumption.

Step 3: Consumption as a function of n_i. Substituting into the budget constraint:

$$c_i + \phi_i n_i + \theta(1-\nu_i)\, c_i = w \qquad \Longrightarrow \qquad c_i\big[1 + \theta(1-\nu_i)\big] = w - \phi_i n_i$$

Hence $c_i = (w - \phi_i n_i)/[1 + \theta(1-\nu_i)]$.

Step 4: FOC for quantity. Differentiating with respect to n_i:

$$\frac{\theta\nu_i + \psi e_i}{n_i} = \frac{\phi_i + \gamma h_i}{c_i}$$

Substituting $\gamma h_i = \theta(1-\nu_i)c_i/n_i$ into the right-hand side:

$$\frac{\theta\nu_i + \psi e_i}{n_i} = \frac{\phi_i}{c_i} + \frac{\theta(1-\nu_i)}{n_i}$$

Rearranging with $A_i \equiv \theta(2\nu_i - 1) + \psi e_i$:

$$\frac{A_i}{n_i} = \frac{\phi_i}{c_i} = \frac{\phi_i\big[1 + \theta(1-\nu_i)\big]}{w - \phi_i n_i}$$

Step 5: Solve for n_i. Cross-multiplying:

$$A_i(w - \phi_i n_i) = \phi_i\big[1 + \theta(1-\nu_i)\big]\, n_i$$

Collecting terms with $D_i \equiv A_i + 1 + \theta(1-\nu_i) = \theta\nu_i + 1 + \psi e_i$:

$$n_i^* = \frac{A_i \cdot w}{\phi_i D_i}$$

Back-substitution yields $c_i^* = w/D_i$ and $h_i^* = \theta(1-\nu_i)\phi_i/(\gamma A_i)$. □

A.2 Proof of Proposition 3.1 (Interior Optimum and Negative Definite Hessian)

Part 1: Interiority. From the closed-form solution (8), $n_i^* = A_i w/(\phi_i D_i)$. Since $A_i > 0$ (by hypothesis), $w > 0$, $\phi_i > 0$, and $D_i = \theta\nu_i + 1 + \psi e_i > 0$, we have $n_i^* > 0$. Similarly, $c_i^* = w/D_i > 0$ and $h_i^* = \theta(1-\nu_i)\phi_i/(\gamma A_i) > 0$ since $\nu_i \in (0,1)$, $\gamma > 0$, and $A_i > 0$.

Part 2: Second-order conditions. Let $\tilde{U}(n_i, h_i) \equiv U_i\big|_{c_i = w - \phi_i n_i - \gamma n_i h_i}$ denote the reduced-form utility. The Hessian entries are:

$$H_{11} \equiv \frac{\partial^2 \tilde{U}}{\partial n_i^2} = -\frac{(\phi_i + \gamma h_i)^2}{c_i^2} - \frac{\theta\nu_i + \psi e_i}{n_i^2} < 0$$

$$H_{22} \equiv \frac{\partial^2 \tilde{U}}{\partial h_i^2} = -\frac{(\gamma n_i)^2}{c_i^2} - \frac{\theta(1-\nu_i)}{h_i^2} < 0$$

$$H_{12} \equiv \frac{\partial^2 \tilde{U}}{\partial n_i \partial h_i} = -\frac{\gamma n_i(\phi_i + \gamma h_i)}{c_i^2} - \frac{\gamma}{c_i}$$

Both diagonal entries are sums of strictly negative terms. To verify the determinant condition $H_{11}H_{22} - H_{12}^2 > 0$, substitute the closed-form solutions. Using the FOC relations $(\phi_i + \gamma h_i)/c_i = (D_i - 1)/n_i$ and $\gamma n_i/c_i = \theta(1-\nu_i)/h_i$, together with the quality-spending identity $\gamma n_i h_i = \theta(1-\nu_i)c_i$, one obtains after algebraic simplification:

$$H_{11}H_{22} - H_{12}^2 = \frac{\gamma^2 D_i^3 A_i}{\theta(1-\nu_i)\, w^2}$$

This expression is strictly positive whenever $A_i > 0$ and $\nu_i \in (0,1)$. Combined with $H_{11} < 0$, the Hessian is negative definite at the interior optimum. □

A.3 Proof of Proposition 3.2 (Existence via Brouwer)

Proof. Under $\nu_i > \frac{1}{2}$ for both groups, the condition $A_i = \theta(2\nu_i - 1) + \psi e_i > 0$ holds for all exposure levels $E_i \geq 0$ (since $\theta(2\nu_i - 1) > 0$ and $\psi e_i \geq 0$). The map $F : \mathbb{R}_+^2 \to \mathbb{R}_+^2$ defined by $F_i(\mathbf{N}) = n_i^*(e(E_i(\mathbf{N})))$ is therefore well defined and continuous.

Self-mapping. Consider the compact rectangle $\mathcal{R} = [0,\, w/\phi_H] \times [0,\, w/\phi_L]$. For any $\mathbf{N} \in \mathcal{R}$:

1. $F_i(\mathbf{N}) > 0$: since $A_i > 0$ for all $E_i \geq 0$, the closed-form gives $n_i^* = A_i w/(\phi_i D_i) > 0$, including at $\mathbf{N} = \mathbf{0}$.
2. $F_i(\mathbf{N}) < w/\phi_i$: since $D_i = A_i + 1 + \theta(1-\nu_i) > A_i$, we have $n_i^* = A_i w/(\phi_i D_i) < w/\phi_i$.

Thus F maps $\mathcal{R}$ into its interior. By Brouwer's fixed-point theorem, F has a fixed point in $\mathcal{R}$, and the fixed point satisfies $N_i^* > 0$ for both groups.

Uniqueness. At calibrated parameters, uniqueness is verified numerically via grid search over $[0.01, 8]^2$ with an 80×80 grid (6,400 candidate cells). Each cell is used as an initial condition for a nonlinear solver (`fsolve` with the steady-state residual $n_i^*(e(E_i(\mathbf{N}))) - N_i$). All converged solutions agree to machine precision, confirming a single fixed point. The spectral radius of the Jacobian at the steady state is 0.123, well below unity, confirming local asymptotic stability and ruling out nearby alternative attractors. □

A.4 Proof of Proposition 4.2 (Positive Subsidy)

Proof. From (25), $\tau_i/c_i = \text{MEB}_i/\lambda_i$. Since $\text{MEB}_i > 0$ (Proposition 3.4), $\lambda_i > 0$, and $c_i > 0$, it follows that $\tau_i^* > 0$. For the comparative static, write $\tau_i^* = w\psi M_i/(\lambda_i D_i)$ where $M_i > 0$ collects the non-ψ factors and $D_i = \theta\nu_i + 1 + \psi e_i$. Then $\partial\tau_i^*/\partial\psi = wM_i(\theta\nu_i + 1)/(\lambda_i D_i^2) > 0$: the linear increase in ψ through the MEB dominates the concave decrease in c_i. □

A.5 Proof of Proposition 4.4 (Targeting by Externality Rank)

Proof. From (30), $\text{MW}_i > \text{MW}_j$ iff $(1/c_i + \text{MEB}_i/(\lambda_i\phi_i)) > (1/c_j + \text{MEB}_j/(\lambda_j\phi_j))$, with $(1+\mu)^{-1}$ cancelling. When $c_i \approx c_j$ (as at calibration: $c_H = c_L = 0.711$), the private-transfer terms cancel, yielding the simplified rule $\text{MEB}_i/(\lambda_i\phi_i) > \text{MEB}_j/(\lambda_j\phi_j)$. When consumption differs, the full comparison applies: the private-transfer component favors the group with lower c_i, while the externality component favors the group with higher externality per dollar. □

A.6 Proof of Proposition 4.5 (Welfare Loss from Uniformity)

Proof. At the group-specific optimum, $\partial W/\partial\tau_i = 0$ for each i. The uniform constraint $\tau_H = \tau_L = \tau$ yields the FOC $\sum_i \partial W/\partial\tau_i = 0$. By the envelope theorem, the welfare loss is second-order:

$$W(\tau,\tau) - W(\tau_H^*, \tau_L^*) \approx \tfrac{1}{2}(\tau - \tau_H^*)^2 W_{HH} + (\tau - \tau_H^*)(\tau - \tau_L^*)W_{HL} + \tfrac{1}{2}(\tau - \tau_L^*)^2 W_{LL}$$

Since $\tau_H^* \neq \tau_L^*$ whenever the externality-per-dollar ratios differ across groups, the loss is strictly positive whenever $\text{MEB}_H/(\lambda_H\phi_H) \neq \text{MEB}_L/(\lambda_L\phi_L)$. □

A.7 Proof of Proposition 4.3 (MCPF Threshold)

Proof. Rearranging (27): a positive subsidy improves welfare iff $\mu < \text{MEB}_i/(\lambda_i\phi_i) + 1/c_i - 1 \equiv \mu_i^*$. The decomposition $\mu_i^* = \text{MEB}_i/(\lambda_i\phi_i) + (1/c_i - 1)$ separates the externality component from the generic transfer component ($1/c_i - 1 > 0$ when $c_i < 1$). At calibration: $\mu_H^* = 0.295 + 0.407 \approx 0.70$, $\mu_L^* = 0.164 + 0.407 \approx 0.57$, both above standard MCPF estimates. □

A.8 Uniqueness Verification: Numerical Method

Uniqueness of the steady state is verified by grid search. An 80×80 grid over $[0.01, 8.0]^2$ in (N_H, N_L) space provides 6,400 initial conditions for the nonlinear solver (`fsolve`, residual in N-space with softplus floor $\epsilon = 0.01$). At baseline parameters, all 6,400 runs converge to the same fixed point: $(3.000, 2.300)$ pre-shock, $(2.501, 1.802)$ post-shock. The Jacobian has spectral radius 0.123 (eigenvalues 0.071, 0.123), well inside the unit circle. The procedure is repeated across all sensitivity-analysis parameter combinations ($\psi \in [0.05, 0.50]$, $\rho \in [0.5, 2.0]$, $\delta \in [0.3, 0.9]$); uniqueness holds in all cases, with the spectral radius remaining well inside the unit circle (maximum 0.28 at $\psi = 0.50$).

B Additional Quantitative Results

This appendix collects supplementary figures and tables referenced in the main text.

B.1 Supplementary Figures

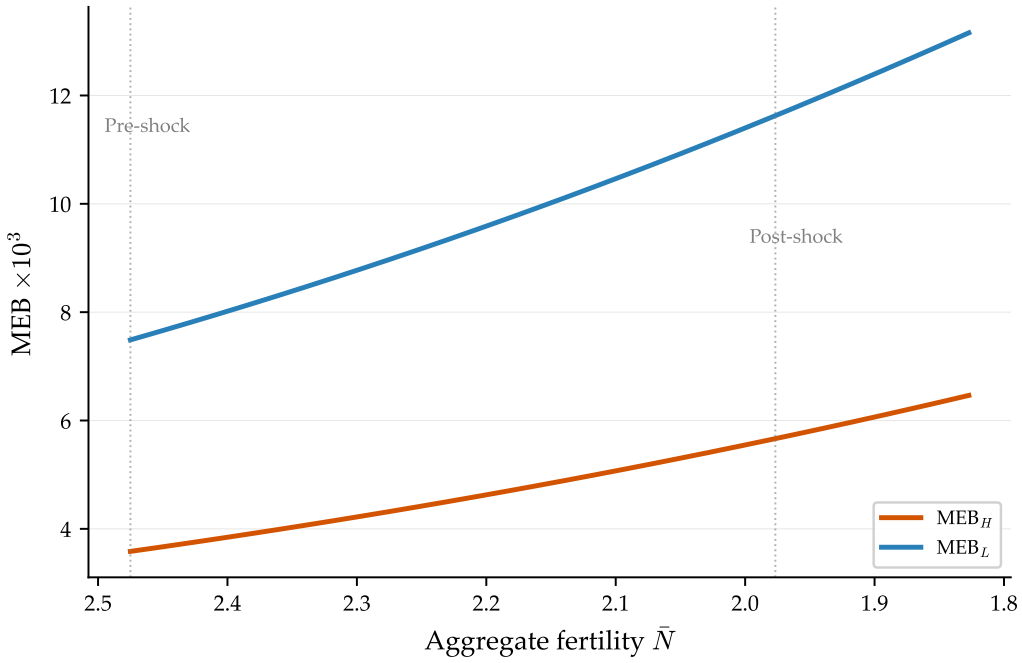

Figure B.1. *Marginal external benefit vs. aggregate fertility. The* MEB *for both groups is plotted against $\bar{N}$ as the cost shock increases from zero to the full calibrated level. Two opposing forces are visible: the empathy sensitivity $e'(E)$ rises as fertility declines (concavity of the empathy function), while the empathy value $\psi \ln(n_j)$ falls. At the calibrated post-shock steady state, the sensitivity channel dominates and the* MEB *is larger than at the pre-shock equilibrium.*

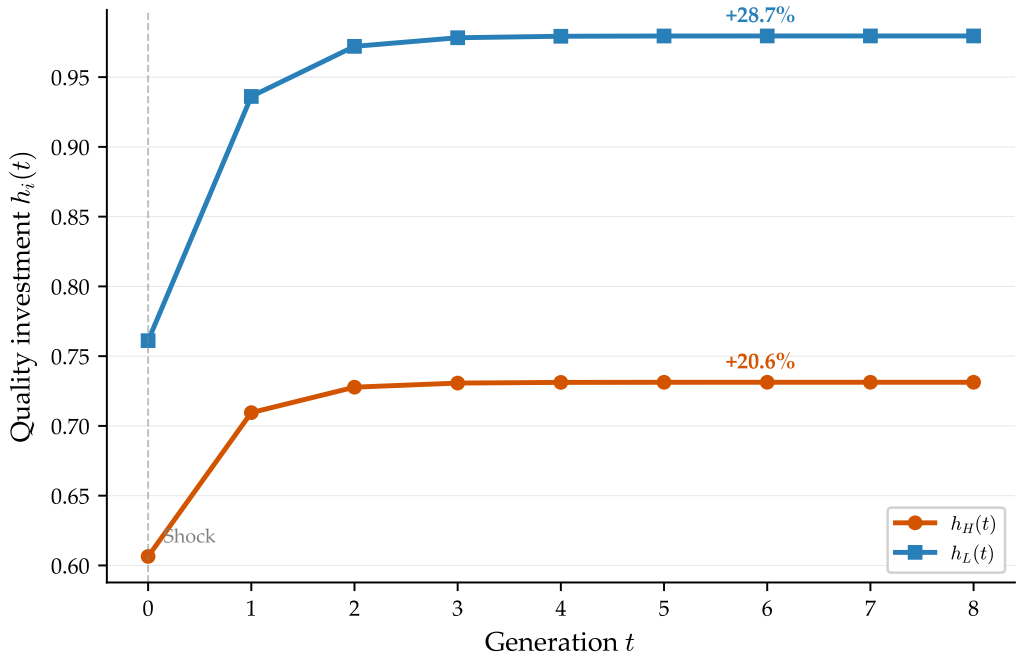

Figure B.2. *Quality response during the transition. Both groups increase quality investment h_i as fertility declines after the cost shock, reflecting the quantity-quality trade-off. Group L's quality rises more steeply (h_L: 0.761 $\rightarrow$ 0.980) than Group H's (h_H: 0.606 $\rightarrow$ 0.731), consistent with the larger cost shock to Group L.*

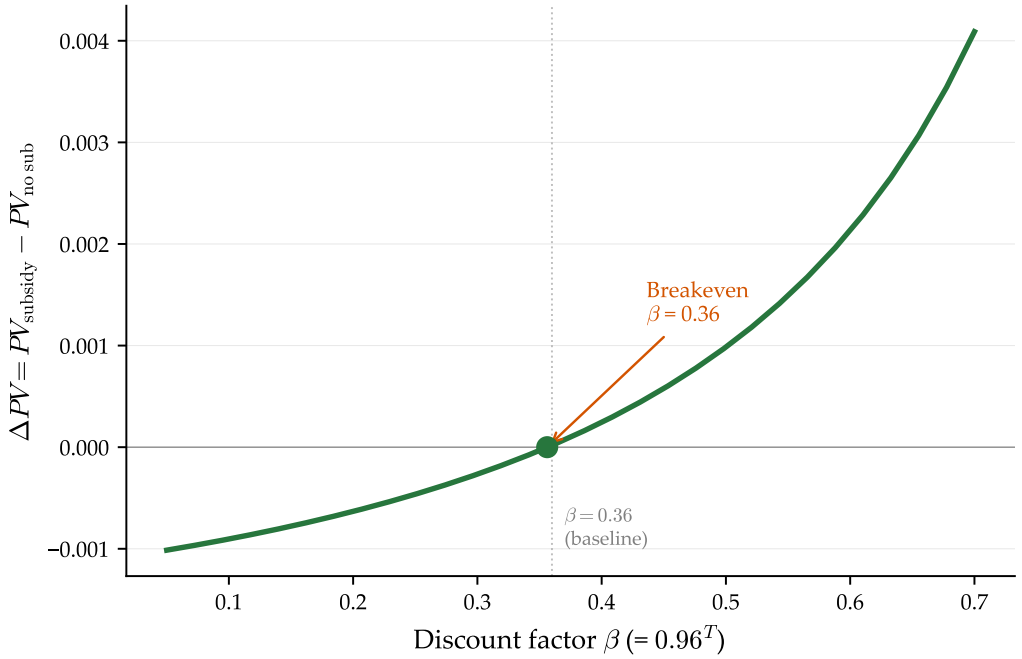

Figure B.3. *Transition-path present value as a function of the discount factor β. The PV of the GE-optimal subsidy (group-specific rates) is evaluated over the transition from pre-shock to post-shock steady state. At $\beta = 0.36$ (baseline: 0.96^{25}), the PV is approximately zero (+0.00002), reflecting the near-exact offset between the front-loaded lump-sum tax and the accumulated steady-state welfare gain. The PV rises with β, turning clearly positive for $\beta > 0.5$.*

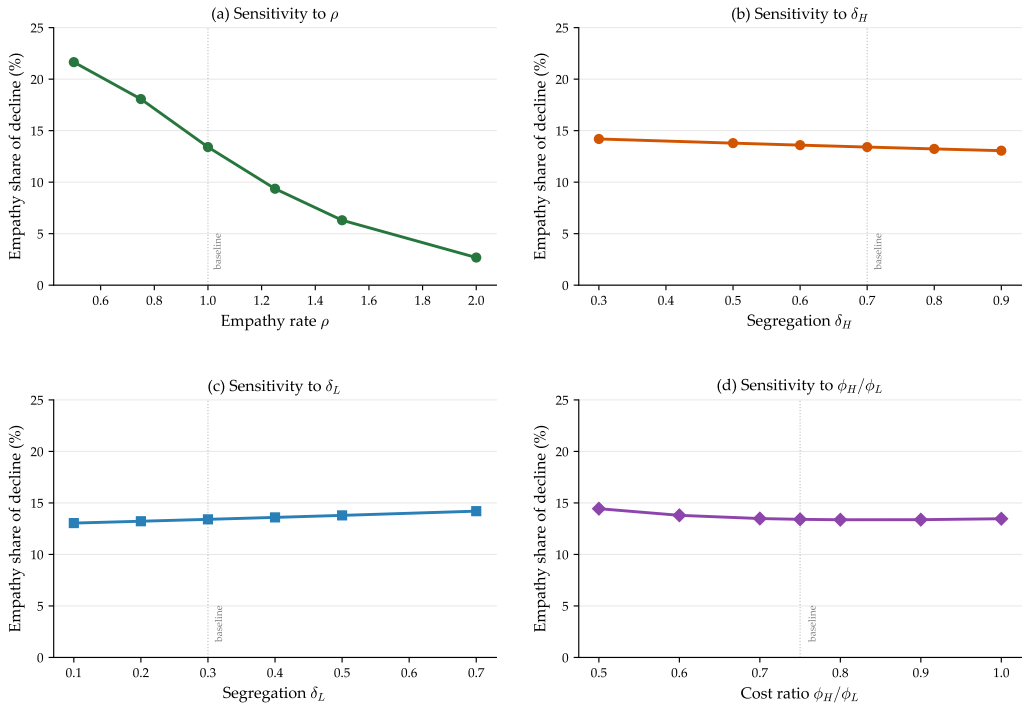

Figure B.4. *Robustness dashboard. The four panels show the sensitivity of the empathy share to structural parameters. Panel (a): empathy rate ρ. Panel (b): segregation of Group H (δ_H). Panel (c): segregation of Group L (δ_L). Panel (d): cost ratio ϕ_H/ϕ_L. The empathy share is most sensitive to ρ and largely insensitive to δ and the cost ratio. Baseline values are marked in each panel.*

B.2 Supplementary Tables

Table B.1. *Transition-Path Present Value Under Alternative Subsidy Scenarios*

Scenario	τ_H	τ_L	ΔW_{SS}	PV ($\beta = 0.36$)
No subsidy (baseline)	0	0	—	—
GE-optimal, group-specific	0.00952	0.00760	+0.00219	+0.00002
GE-optimal, uniform	0.00845	0.00845	+0.00216	+0.00000
1st-order Pigouvian, group-specific	0.01257	0.00935	+0.00197	−0.00073
1st-order Pigouvian, uniform	0.01015	0.01015	+0.00206	−0.00048
GE-optimal, group-spec (post-shock)	0.01300	0.00956	+0.00265	−0.00087
GE-optimal, uniform (post-shock)	0.01124	0.01124	+0.00258	−0.00089
GE-optimal, delayed $t = 3$	0.01300	0.00956	+0.00265	−0.00003

Notes. ΔW_{SS} is the steady-state welfare gain relative to no subsidy. PV is the transition-path present value using generational discount factor $\beta = 0.36$ ($= 0.96^{25}$). Negative PV indicates that the front-loaded fiscal cost exceeds the discounted welfare benefit along the transition path, even though the steady-state gain is positive. The 1st-order Pigouvian rates overshoot the GE optimum by 23–32%, producing negative transition-path PV.

Table B.2. *Combined Structural Robustness*

ψ	Empathy share	τ_H^*	τ_L^*	ΔW_{SS}	μ_H^*	μ_L^*
0.05	3.4%	0.00303	0.00237	0.00015	0.40	0.37
0.10	6.8%	0.00561	0.00438	0.00058	0.50	0.44
0.20	13.4%	0.00952	0.00760	0.00219	0.70	0.57
0.30	19.9%	0.01223	0.01003	0.00466	0.90	0.70
0.40	26.5%	0.01417	0.01188	0.00782	1.10	0.84
0.50	33.3%	0.01559	0.01331	0.01154	1.30	0.97

Notes. GE-optimal group-specific subsidy rates and associated welfare metrics across the full range of ψ. At each value of ψ, the quantity weights ν_H and ν_L are recalibrated to match the same pre-shock fertility targets ($n_H = 3.0$, $n_L = 2.3$). ΔW_{SS} is the steady-state welfare gain from the GE-optimal subsidy. The MCPF thresholds μ_i^* exceed the lower end of the empirical range ($\mu = 0.1$) for all $\psi \geq 0.05$; they exceed the upper end ($\mu = 0.5$) for both groups when $\psi \geq 0.20$. At $\psi = 0.05$, both groups' thresholds fall below 0.50; at $\psi = 0.10$, Group L's threshold (0.44) remains below 0.50 while Group H's (0.50) is borderline. The subsidy may not survive the highest MCPF estimates at these low values of ψ.