Generosity, Legibility, and Fertility Inertia

HongSeok Kim

Department of Economics, Iowa State University

Ames, Iowa, USA

hongseok@iastate.edu

Abstract

I study how public information and credible transfers affect fertility through different channels. The model extends dynastic altruism with risk-sensitive preferences over total descendant utility and maxmin ambiguity about signal informativeness. Building on information inertia, a local normal representation yields an interval of adverse news over which welfare falls while fertility, investment per child, and consumption remain locally unchanged. Insensitivity to news nevertheless coexists with sensitivity to a credible subsidy. In a closed-form policy specialization, an ambiguous transfer is equivalent to a certain transfer scaled by its lowest admissible delivery rate. Greater effective support raises fertility but also shifts and widens the inertia interval, whereas narrowing ambiguity about signal informativeness reduces its width for a fixed family plan. The analysis distinguishes statutory generosity, minimum credible benefit delivery, and uncertainty about public information. Extensions consider smooth ambiguity and integer fertility.

Keywords: Fertility; family policy; legibility; ambiguity aversion; information quality

JEL classification: J13, J18, D81

1 Introduction

Having a child commits a household to a future it cannot fully predict. Public news about labor markets, housing, or family policy can inform that decision without telling prospective parents how closely the announcement applies to their own children. When does information change how parents evaluate their future without changing their family plans, and what does that distinction imply for family policy? This paper studies fertility responsiveness through three margins: public information, statutory generosity, and credible benefit delivery. Parents can leave their family plans unchanged over a range of adverse signals while still responding to a credible transfer. I use *legibility* to describe the household's ability to determine, before having a child, what benefit it can count on receiving; the policy model isolates its minimum credible delivery rate.

Family-policy experience motivates this distinction. South Korea's total fertility rate fell from 1.13 in 2006 to 0.72 in 2023 despite a substantial expansion of policies intended to support births (Statistics Korea, 2024).[1] This aggregate record does not identify the effect of policy: fertility could have fallen further without the interventions. It does, however, motivate a distinction between public commitments and the support a household expects to be available over its planning horizon.

Program evaluations give the distinction more concrete content. Studies of birth bonuses and child subsidies document fertility responses in Quebec, Spain, Israel, and South Korea (Milligan, 2005; González, 2013; Cohen et al., 2013; Kim, 2024). For Korean district bonuses, Kim (2024) estimates completed-fertility elasticities of 0.03–0.06; benefits were locally financed and renewed annually. Evaluations of Germany's earnings-related parental benefits and childcare expansion also find fertility effects (Raute, 2019; Bauernschuster et al., 2016). These studies concern different populations and policy margins and do not identify a common ranking of births per dollar or a causal effect of legibility. They motivate asking how the timing, eligibility, and delivery of a benefit enter a household's valuation alongside its announced amount.

Research on subjective uncertainty provides a complementary motivation. Vignoli et al. (2020) and Guetto et al. (2022) emphasize personal expectations and narratives about the future in fertility intentions. Among young Korean adults, Kwag et al. (2024) documents an association between institutional trust and fertility intentions. Such findings motivate uncertainty about one's own future circumstances; they do not establish maxmin preferences or ambiguity about signal informativeness. Likewise, a decline in births after a recession and a small response to a policy announcement do not identify an asymmetric response to comparable news. The asymmetry and the interval of inaction derived below are predictions of the theory.

I extend the Barro–Becker dynastic structure with risk-sensitive preferences and ambiguity about signal informativeness. Parents choose consumption, fertility, and investment per child. Children share exposure to a common future environment, and the parent evaluates their aggregate continuation value through a risk-sensitive certainty equivalent. After observing a public signal, she considers a set of possible values for its informativeness and evaluates each candidate family plan under the least favorable value in that set. The adverse interpretation is chosen separately for each plan. It is

[1] The cumulative spending figure reported in early 2024 was roughly 380 trillion won across the government's Basic Plans on Low Fertility and Aged Society (Al Jazeera, 2024). This is a broad policy-budget measure, not a measure of transfers received by prospective parents.

therefore an endogenous part of the decision problem rather than a fixed pessimistic belief. Dynastic altruism supplies the link between parental choices and descendants' utility; the risk-sensitive aggregation of that utility is an additional preference assumption.

The mechanism rests on two effects of informative bad news. It lowers the conditional mean of the children's continuation value, but it also resolves some of their common risk. Because additional children increase the dynasty's exposure to that risk, the balance between these effects depends on the family plan. Treating adverse news as more informative makes its content more damaging while leaving less risk unresolved. Treating it as less informative softens the mean loss but increases the risk penalty. Over an intermediate range of adverse signals, the least-favorable balance lies strictly inside the informativeness set and adjusts with both the news and the proposed plan.

Building on the information-inertia mechanism of Illeditsch et al. (2021), the first result characterizes a separation between welfare and the entire family allocation. Allowing investment per child to adjust need not restore responsiveness to news. Under a local normal representation of continuation value, minimization over an interior informativeness parameter leaves a signal-dependent term that is independent of fertility and investment. Worse news reduces value but leaves the marginal incentives unchanged. Under the stated local optimality conditions, fertility, investment per child, and consumption are therefore constant throughout an interval of moderately adverse signals (Proposition 4). I call this *fertility inertia*. The parent pays attention and revises her evaluation; the adjustment occurs in the worst-case interpretation rather than in the family plan. This result is local and holds current resources and child costs fixed. The normal representation is a tractable reduction motivated by the recursive framework, not a consequence of its bounded-value existence theorem.

A closed-form specialization makes the boundaries of inertia explicit. Favorable signals are evaluated at the lowest admissible informativeness; sufficiently adverse signals are evaluated at the highest. Away from resource and nonnegativity constraints, fertility consequently has three regions: a steeper response to sufficiently bad news, a flat interval, and a shallower response to good news (Corollary 2). A wider informativeness set expands the interval for a fixed signal-free family plan. Under additional regularity conditions, smooth ambiguity yields a locally attenuated response that approaches the maxmin result as ambiguity aversion grows. With integer fertility, the value difference between two positive family sizes is constant wherever both have interior worst-case beliefs. The comparison with childlessness is different because a childless plan has no exposure to continuation risk.

The policy application distinguishes responsiveness to news from responsiveness to material incentives. Inside the inertia interval, fertility is locally insensitive to the public signal but responds positively to an increase in the subsidy when its minimum credible delivery rate is positive. Thus a flat response to news does not imply an inactive fertility margin. Suppose a statutory per-child subsidy τ reaches a household at an ambiguous pass-through rate $\alpha \in [\underline{\alpha}, \overline{\alpha}]$. In the policy specialization, the household evaluates it at $\underline{\alpha}\tau$, so the program is equivalent to a certain subsidy of that amount (Proposition 5). For $\underline{\alpha} > 0$, matching a certain entitlement requires scaling the ambiguous promise by $1/\underline{\alpha}$. Realized expenditure can nevertheless reflect a larger pass-through. This is a theoretical wedge between the transfer that guides choice and the transfer ultimately paid, rather than an estimate of the effectiveness of any particular program. Greater effective support also raises the parent's exposure and shifts and widens the inertia interval in the closed-form specialization. This width result does not imply lower responsiveness at every signal or a higher probability of inertia. Raising the minimum credible pass-through strengthens the subsidy's fertility incentive, but reducing inertia additionally requires a sufficiently large narrowing of ambiguity about signal informativeness. Benefit delivery and the interpretation of news are distinct margins.

Related literature The paper connects the quantity–quality approach to fertility (Becker and Lewis, 1973) and its dynamic and dynastic formulations (Razin and Ben-Zion, 1975; Becker and Barro, 1988; Barro and Becker, 1989) to ambiguity about information quality. Subsequent work studies fertility alongside growth, human capital, and inequality (Galor and Weil, 2000; de la Croix and Doepke, 2003; Doepke and Zilibotti, 2005; Manuelli and Seshadri, 2009), and under income, employment, and other risks (Ranjan, 1999; Sommer, 2016; Doepke et al., 2023). Here the additional uncertainty concerns the probability model linking public news to a dynasty's continuation value. The contribution is a characterization of joint fertility and investment inertia and its implications for the distinct effects of information and credible transfers. The recursive framework extends the treatment of uncertainty; it does not obtain risk-sensitive preferences by reexpressing the original Barro–Becker objective.

The decision-theoretic ingredients come from maxmin utility (Gilboa and Schmeidler, 1989), recursive ambiguity (Epstein and Schneider, 2003), risk-sensitive and robust control (Hansen and Sargent, 1995, 2001, 2008), and smooth ambiguity (Klibanoff et al., 2005, 2009). The closest mechanism is information-quality ambiguity in financial markets. Epstein and Schneider (2008) derive asymmetric responses to news; related analyses include Caskey (2009) and Ju and Miao (2012). Most directly, Illeditsch et al. (2021) show how risk and ambiguity jointly generate information inertia in portfolio choice. This paper carries that mechanism into fertility: dynastic continuation value provides the return, family size scales common-risk exposure, and investment changes the return's mean and risk. The fertility application also permits an explicit distinction between statutory generosity and ambiguous benefit delivery. Inattention, adjustment costs, and social norms remain alternative or complementary explanations for weak fertility responses; the results here isolate the implications of plan-dependent interpretations of public information.

Outline Section 2 presents the recursive environment and its optimality conditions. Section 3 derives fertility inertia in the local normal model, develops the policy application, and discusses smooth ambiguity and integer fertility. Section 4 concludes. Proofs and extensions are collected in the appendices.

2 Dynastic Fertility with Risk Sensitivity and Ambiguous Signals

This section extends the Barro–Becker dynastic structure with a risk-sensitive evaluation of descendant utility and ambiguous public information. I specify the household's preferences and information, establish existence of a recursive solution, and derive the optimality conditions. Section 3 uses a local normal representation to characterize the inertia mechanism and its policy implications.

2.1 Environment

The state. Time is discrete and each period corresponds to one generation. A parent enters the period with state

$$(x_t, s_t) \in \mathcal{X} \times \mathcal{S}.$$

The vector x_t contains the payoff-relevant household and aggregate state: assets, human capital, age, economic conditions, and policy variables. The scalar s_t is public news about the environment in which the next generation would grow up. It may concern housing costs, childcare, labor markets, or family policy. Larger values of s_t represent better news. Current resources are $m_t \equiv y_t + a_t$, with income y_t and assets a_t determined by x_t.

The controls. The controls are consumption $c_t \geq 0$, fertility $n_t \geq 0$, and a transfer or investment $z_t \geq 0$ for each child. The last control is the quality margin in the familiar quantity–quality tradeoff. I treat fertility as continuous; a fractional value can be interpreted as the intensity or timing of an additional-birth plan. The budget constraint is

$$c_t + n_t(q_t + z_t) \leq m_t, \tag{1}$$

where $q_t > 0$ is the direct cost of a child. Total resources devoted to each child are therefore $q_t + z_t$.

The transition and the ambiguous parameter. The next generation's state is generated by

$$(x_{t+1}, s_{t+1}) = G(x_t, s_t, z_t, \varepsilon_{t+1}; \rho), \tag{2}$$

The innovation ε_{t+1} is realized after the parent chooses and has a distribution independent of those choices. The parameter ρ indexes the probability models the parent regards as possible. In particular, it measures how strongly the public signal predicts the next generation's continuation value. The parent knows that the signal is relevant but cannot pin down its informativeness:

$$\rho \in [\underline{\rho}, \overline{\rho}], \qquad 0 < \underline{\rho} < \overline{\rho} < 1, \tag{3}$$

and assigns no prior over this interval. Aggregate news about housing or childcare is plainly relevant to parents as a group, yet a household may not know how well it predicts the circumstances of its own descendants. The uncertainty over ρ is Knightian; it is separate from ordinary risk over ε_{t+1}. I write expectations under model ρ and investment z_t as $\mathbb{E}_{\rho,z_t}[\,\cdot \mid x_t, s_t]$, suppressing subscripts when they are clear.

Timing. The timing is straightforward. The parent first sees s_t, then chooses (c_t, n_t, z_t), and finally observes ε_{t+1}. Thus the choice is made with risk over the innovation and ambiguity over its signal-loading parameter ρ.

The dynastic objective. Let

$$V' \equiv V(x_{t+1}, s_{t+1})$$

denote a child's lifetime continuation value, including her own future fertility choices. The same value function applies in every generation, so the return to a child is endogenous rather than a primitive payoff. Fertility scales that continuation value through the Barro–Becker factor $\Phi(n_t)$, where

$$\Phi(0) = 0, \qquad \Phi'(n) > 0, \qquad \Phi''(n) \le 0.$$

With no children, the dynasty receives no continuation payoff. Additional children increase dynastic weight, though weakly at a decreasing rate. A standard specification is $\Phi(n) = n^{1-\eta}/(1-\eta)$ for $\eta \in [0, 1)$; the case $\eta = 0$ gives $\Phi(n) = n$. The results do not rely on this particular form.

2.2 Risk-sensitive continuation value

Relation to Barro–Becker altruism. The dynastic structure follows Becker and Barro (1988); Barro and Becker (1989): parents value their own consumption and the utility of their children, who face the same type of decision problem. In a conventional expected-utility extension to uncertainty, the continuation term would be $\beta\Phi(n_t)\mathbb{E}[V' \mid x_t, s_t]$. This expression is linear in expected continuation utility; it need not be linear in descendants' consumption. Concavity of utility can already make consumption risk costly. The present model adds a separate aversion to risk in total dynastic continuation utility. This changes preferences over uncertain family outcomes while preserving the altruistic recursion under certainty.

When the parent chooses, V' is still random. I summarize that risk with an exponential certainty equivalent. For any continuation payoff Y, define

$$\mathcal{C}(Y \mid x_t, s_t) = \mathcal{C}_{\gamma,\rho,z_t}(Y \mid x_t, s_t) \equiv -\frac{1}{\gamma}\log \mathbb{E}\left[\exp(-\gamma Y) \mid x_t, s_t\right], \qquad \gamma > 0. \tag{4}$$

This is the sure payoff that a CARA decision maker with absolute risk aversion γ considers equivalent to Y. Greater γ places a larger penalty on risk, while the limit as $\gamma \to 0$ is the ordinary expectation. The operator is standard in risk-sensitive control (Hansen and Sargent, 1995; Tallarini, 2000). For normal payoffs it has the exact mean–variance form stated next.

Lemma 1 (CARA certainty equivalent)**.** *$\mathcal{C}_{\gamma,\rho,z_t}(Y \mid x_t, s_t)$ is the sure payoff a CARA agent with utility $U(Y) = -\exp(-\gamma Y)$ regards as equivalent to Y under model ρ. If $Y \mid x_t, s_t, \rho \sim N(\mu_Y, \sigma_Y^2)$, the formula is exact and linear–quadratic,*

$$\mathcal{C}(Y \mid x_t, s_t) = \mu_Y - \frac{\gamma}{2}\sigma_Y^2.$$

For general Y whose cumulant-generating function $\kappa(t) = \log \mathbb{E}e^{tY}$ is finite and twice differentiable in a neighborhood of $t = 0$, a small-γ expansion gives $\mathcal{C}(Y) = \mathbb{E}[Y] - \frac{\gamma}{2}\operatorname{Var}(Y) + o(\gamma)$, so $\lim_{\gamma\to 0}\mathcal{C}(Y) = \mathbb{E}[Y]$.

Appendix A.1 gives the derivation. Under normality, risk subtracts $\gamma/2$ times the variance from the mean. The same expression is the leading approximation for a general payoff when γ is small.

Risk is evaluated after aggregating descendant utility. The location of fertility inside the certainty equivalent is substantive. For $\Phi(n) > 0$,

$$\mathcal{C}_\gamma[\Phi(n)V'] = \Phi(n)\mathcal{C}_{\gamma\Phi(n)}[V'] \neq \Phi(n)\mathcal{C}_\gamma[V'] \quad \text{in general.}$$

Under conditional normality, the model's aggregate evaluation is

$$\mathcal{C}_\gamma[\Phi(n)V'] = \Phi(n)\mathbb{E}[V'] - \frac{\gamma}{2}\Phi(n)^2 \operatorname{Var}(V'),$$

whereas applying the certainty equivalent to each child's value before scaling would give a risk penalty proportional to $\Phi(n)$. The squared exposure in the present specification comes from evaluating the common shock after aggregating descendant utility. It is the preference assumption that allows the least-favorable informativeness to vary with fertility in the normal model. The mean–variance formula is exact under the exponential certainty equivalent and normality; it is not an approximation to the original Barro–Becker objective.

This distinction matters for the mechanism. Under the normal representation used below, ordinary expected continuation utility is $\Phi(n)[\mu(z) + \sigma(z)\rho s]$. For a positive family plan, its minimizer over ρ is the lower endpoint when $s > 0$ and the upper endpoint when $s < 0$. There is no residual-variance term to balance the mean effect and no interior worst-case adjustment of the type used in the inertia result. This comparison holds for the maintained normal representation; it does not rule out other mechanisms in stochastic models with nonlinear continuation values.

The recursive problem. The payoff to fertility is the scaled continuation value

$$Y_{t+1}(n_t, z_t) = \Phi(n_t)V(x_{t+1}, s_{t+1}). \tag{5}$$

The parent evaluates this payoff in two stages. She first computes its certainty equivalent under a particular model and then takes the least favorable value over the admissible models. This is the maxmin ordering of Gilboa and Schmeidler (1989): each action is judged under its own worst plausible probability law. The recursive problem is

$$V(x_t, s_t) = \max_{c_t, n_t, z_t} \left\{ u(c_t) + \beta \min_{\rho \in [\underline{\rho}, \overline{\rho}]} \mathcal{C}_{\gamma,\rho,z_t}\Big[\Phi(n_t)V(x_{t+1}, s_{t+1}) \mid x_t, s_t\Big] \right\} \tag{6}$$

subject to (1). Flow utility satisfies $u' > 0$ and $u'' < 0$, and $\beta \in (0, 1)$. The problem retains a dynastic fixed-point structure: the value of a child is the next generation's value function. The nonlinear evaluation changes preferences and generally changes the fixed point itself. Under certainty the continuation term reduces to $\beta\Phi(n_t)V'$. The order of the two operators is important. Nature selects ρ separately for each proposed (n_t, z_t), allowing the least-favorable reading of the news to change with the family plan.

Interpretation The certainty equivalent captures unavoidable risk. The minimization over ρ captures uncertainty about how to read the signal. These two sources of caution interact differently with different plans. A large family has more exposure to the common future and therefore more to lose from unresolved risk; its adverse reading of moderately bad news places less weight on the signal. For a smaller family, the mean effect matters more, so the adverse reading is more informative. If the pessimistic belief were fixed, the two effects would balance at only one plan. Here the belief moves with the plan, allowing the balance to persist over an interval and creating a flat policy region.

2.3 Recursive problem and solution concept

The value function appears on both sides of (6), so existence and uniqueness are not automatic. Define the Bellman operator T by

$$(TV)(x,s) = \max_{c,n,z} \left\{ u(c) + \beta \min_{\rho\in[\underline{\rho},\overline{\rho}]} \mathcal{C}_{\gamma,\rho,z}\Big[\Phi(n)V(x',s') \mid x,s\Big] \right\} \tag{7}$$

subject to

$$c + n\Big(q(x,s)+z\Big) \le m(x,s).$$

A solution is a fixed point $V = TV$ together with policy functions

$$c^*(x,s), \qquad n^*(x,s), \qquad z^*(x,s),$$

and a worst-case belief policy

$$\rho^*(x,s) \in [\underline{\rho},\overline{\rho}].$$

Assumption 1 (Regularity). *(i) $\mathcal{X}\times\mathcal{S}$ is a Polish space. (ii) Per-child investment is bounded, $z\in[0,\bar z]$, so that the feasible correspondence $\Gamma(x,s) = \{(c,n,z)\ge 0 : c+n(q(x,s)+z) \le m(x,s),\ z\le \bar z\}$ is nonempty, compact-valued, and continuous. (iii) u is continuous, increasing, concave, and bounded on the graph of Γ, with $\sup_{(x,s),\,a\in\Gamma(x,s)} |u(c)| < \infty$. (iv) Φ is continuous and increasing with $\overline{\Phi} \equiv \sup_{n\in\Gamma_n}\Phi(n) < \infty$, where Γ_n is the (compact) set of feasible fertility levels, and $\beta\overline{\Phi} < 1$. (v) The transition law $P_{\rho,z}(\cdot \mid x,s)$ is weakly continuous (Feller) in (x,s,z) and continuous in ρ. The law of the innovation ε_{t+1} does not depend on the controls, so the subscript z in $\mathbb{E}_{\rho,z}$ records only that the successor state depends on z through G. The operator T acts on the Banach space of bounded continuous functions on $\mathcal{X}\times\mathcal{S}$ equipped with the sup norm.*

Assumption 1 follows the recursive ambiguity framework of Epstein and Schneider (2003). One issue requires care. Repeated period-by-period minimization need not be dynamically consistent unless the set of one-step beliefs is rectangular and updated prior by prior. Treating ρ as one unknown parameter fixed forever would generally fail that requirement. I therefore state the intergenerational ambiguity set explicitly.

Assumption 2 (Generation-by-generation ambiguity)**.** *For each state and action, nature's set of one-step-ahead transition laws is the kernel correspondence*

$$\mathcal{P}(x, s, z) = \left\{ P_{\rho,z}(\cdot \mid x, s) : \rho \in [\underline{\rho}, \overline{\rho}] \right\}.$$

The parent's set of priors over future paths is the rectangular hull of this correspondence, the set of all path measures generated by measurable, possibly history-dependent selections $\rho_t(h_t) \in [\underline{\rho}, \overline{\rho}]$ *from* $\mathcal{P}$ *at every date and history.*

This rectangular structure has a direct economic interpretation. Each generation faces a new housing market, labor market, and policy regime, so the link between that generation's headlines and its private outcomes can change over time. Rectangularity makes the maxmin recursion dynamically consistent and ensures that plan-dependent worst-case beliefs come from the decision problem itself. The CARA layer is not covered by the original Epstein–Schneider theorem; I take the composed recursion (6) as the primitive preference specification, following the spirit of recursive robust control (Hansen and Sargent, 2008). Appendix A.7 discusses the issue and gives the entropy dual.

Remark 1 (No learning about informativeness). Under Assumption 2 the dynasty never learns about ρ. Nature may select a new ρ_t at every date and history, so past signals carry no information about the informativeness of future ones. This is not an incidental omission. Epstein and Schneider (2007) show that when ambiguity attaches to the quality of signals, even long histories need not shrink the set of beliefs. The smooth-ambiguity variant sketched in Section 3.7 is the natural framework for learning extensions, since it accommodates nondegenerate posteriors over ρ.

The condition $\beta\overline{\Phi} < 1$ plays the role ordinarily assigned to $\beta < 1$. Here the effective one-generation discount factor is $\beta\Phi(n)$ because a parent with more children places greater total weight on descendants. Uniformly bounding this product below one keeps the dynastic value finite and supplies the contraction.

Proposition 1 (Existence and characterization)**.** *Under Assumption 1, T is a contraction on the bounded continuous functions with modulus at most $\beta\overline{\Phi}$, so a unique $V = TV$ exists. It is attained by policy functions solving the outer maximum and a worst-case belief $\rho^*(x, s)$ solving the inner minimum.*

Proof. The proof is available in Appendix A.2. □

The additional preference layers do not alter the contraction modulus. Both are nonexpansive in the sup norm, so T inherits the standard Barro–Becker bound $\beta\overline{\Phi}$. Ambiguity and risk sensitivity change the location of the optimum, not the existence of the fixed point.

2.4 Equilibrium characterization

The general first-order conditions retain the Barro–Becker structure. Their distinctive feature is the probability measure used to value future outcomes: it is distorted toward unfavorable states by the parent's own caution.

For a given state (x_t, s_t), define the model-ρ certainty equivalent of a plan,

$$K_\rho(n, z; x_t, s_t) \equiv \mathcal{C}_{\gamma,\rho,z}\Big[\Phi(n)V(x_{t+1}, s_{t+1}) \mid x_t, s_t\Big], \tag{8}$$

and its worst case across models,

$$J(n, z; x_t, s_t) \equiv \min_{\rho\in[\underline{\rho},\overline{\rho}]} K_\rho(n, z; x_t, s_t). \tag{9}$$

Since $u'(c) > 0$, the budget constraint binds at an interior solution, $c_t^* = m_t - n_t^*(q_t + z_t^*)$. Hence the reduced problem is

$$\max_{n,z} \Big\{u\Big(m_t - n(q_t + z)\Big) + \beta J(n, z; x_t, s_t)\Big\}. \tag{10}$$

Let

$$\rho_t^* \in \arg\min_{\rho\in[\underline{\rho},\overline{\rho}]} K_\rho(n_t^*, z_t^*; x_t, s_t)$$

be the worst-case signal-informativeness parameter at the optimum. The associated measure Q_t^* begins with the transition law $P_{\rho_t^*, z_t^*}$ and reweights successor states in proportion to $\exp\{-\gamma\Phi(n_t^*)V(x_{t+1}, s_{t+1})\}$. Low-value states receive more weight when $\gamma\Phi(n_t^*) > 0$. Caution therefore enters twice: ρ_t^* selects the least favorable informativeness, and the CARA operator tilts that law toward poor continuation outcomes. When $\gamma \to 0$, the tilt disappears and Q_t^* returns to $P_{\rho_t^*, z_t^*}$. Appendix A.8 records the Radon–Nikodym derivative and the equivalent entropy representation.

Assume that the transition admits the differentiable shock representation (2), with V differentiable along the relevant directions, and define the marginal effect of child investment on the child's continuation value by

$$
\begin{aligned}
D_z V_{t+1} \equiv \nabla_x V(x_{t+1}, s_{t+1}) \cdot G_z^x(x_t, s_t, z_t, \varepsilon_{t+1}; \rho_t^*) \\
+ V_s(x_{t+1}, s_{t+1})\, G_z^s(x_t, s_t, z_t, \varepsilon_{t+1}; \rho_t^*),
\end{aligned}
\tag{11}
$$

where G^x and G^s are the corresponding components of the transition. Thus $D_z V_{t+1}$ is the state-contingent marginal return to one more unit of investment in each child.

Proposition 2 (Interior first-order conditions)**.** *Assume the transition admits the differentiable shock representation* (2) *and that V is differentiable. Let $D_z V_{t+1}$ be defined by* (11)*. Suppose (c_t^*, n_t^*, z_t^*) is an interior optimum, the inner minimizer ρ_t^* is unique, and differentiation under the expectation is valid. Then*

$$
\rho_t^* \in \arg\min_{\rho \in [\underline{\rho}, \overline{\rho}]} K_\rho(n_t^*, z_t^*; x_t, s_t),
$$

and, letting Q_t^* denote the twisted measure associated with (ρ_t^*, n_t^*, z_t^*), the optimal policies satisfy

$$
c_t^* = m_t - n_t^*(q_t + z_t^*), \tag{12}
$$

$$
u'(c_t^*)(q_t + z_t^*) = \beta \Phi'(n_t^*) \mathbb{E}_{Q_t^*}\left[V(x_{t+1}, s_{t+1}) \mid x_t, s_t \right], \tag{13}
$$

$$
n_t^* u'(c_t^*) = \beta \Phi(n_t^*) \mathbb{E}_{Q_t^*}\left[D_z V_{t+1} \mid x_t, s_t \right]. \tag{14}
$$

If $K_\rho(n_t^*, z_t^*; x_t, s_t)$ is differentiable in ρ, the worst-case belief also satisfies the one-dimensional optimality condition

$$
\begin{cases}
\partial_\rho K_\rho(n_t^*, z_t^*; x_t, s_t)\Big|_{\rho = \rho_t^*} = 0, & \rho_t^* \in (\underline{\rho}, \overline{\rho}), \\[2ex]
\partial_\rho^+ K_\rho(n_t^*, z_t^*; x_t, s_t)\Big|_{\rho = \underline{\rho}} \geq 0, & \rho_t^* = \underline{\rho}, \\[2ex]
\partial_\rho^- K_\rho(n_t^*, z_t^*; x_t, s_t)\Big|_{\rho = \overline{\rho}} \leq 0, & \rho_t^* = \overline{\rho}.
\end{cases}
\tag{15}
$$

Proof. The proof is available in Appendix A.3. □

The equations have the usual Barro–Becker interpretation, but marginal benefits are computed under Q_t^*. Because the inner minimizer is unique, the envelope theorem removes any derivative of the worst-case belief with respect to the plan. This is the familiar maxmin structure; an analogous business-cycle application appears in Ilut and Schneider (2014). In (13), the consumption cost of another child equals that child's discounted dynastic value under the distorted measure, weighted by $\Phi'(n_t^*)$. In (14), an additional unit invested in every child costs n_t^* units today and returns $D_z V_{t+1}$ for each child, weighted by $\Phi(n_t^*)$.

Dividing (13) by (14) eliminates the marginal utility of consumption and gives the quantity–quality condition

$$\frac{q_t + z_t^*}{n_t^*} = \frac{\Phi'(n_t^*)}{\Phi(n_t^*)} \frac{\mathbb{E}_{Q_t^*}[V(x_{t+1}, s_{t+1}) \mid x_t, s_t]}{\mathbb{E}_{Q_t^*}[D_z V_{t+1} \mid x_t, s_t]}. \tag{16}$$

The left side compares the resource cost of another child with the cost of raising investment per child. The right side combines the marginal-to-total dynastic weight with the distorted expected value of a child relative to the distorted return on child investment. It is the quantity–quality condition under ambiguous signal informativeness.

Equation (15) determines the adverse reading of the news. At an interior solution, a marginal change in informativeness leaves the minimized objective unchanged. At the endpoints, the worst case is the least or most informative admissible model. The inertia result arises when the interior condition continues to hold over a range of signals.

The distortion remains present in the ratio (16). Unless continuation value and the return to investment move together across states, pessimism reweights them differently. Ambiguity can therefore change the quantity–quality mix rather than merely shifting fertility up or down.

Remark 2 (Special cases). The model nests the standard benchmarks. As $\gamma \to 0$ the certainty equivalent converges to a conditional expectation (Lemma 1) and the problem becomes a risk-neutral robust Barro–Becker problem, in which $\Phi(n_t)$ factors out of the inner minimization. If instead the ambiguity set is a singleton, $\underline{\rho} = \overline{\rho} = \rho$, the problem is ambiguity-neutral but risk sensitive. Switching off both features recovers the standard Barro–Becker recursion. The explicit Bellman equations for each case are collected in Appendix A.9.

3 Signal Informativeness and Fertility Inertia

Closed-form policies are not available in the full recursion. To make the mechanism visible, I work with a local normal representation of the next generation's value. The reduction keeps the dynastic interpretation while isolating the effect of signal informativeness. It produces the paper's central result: welfare responds to moderately bad news even though fertility, investment, and consumption remain fixed.

3.1 Local normal representation

Around the current state, suppose

$$V(x_{t+1}, s_{t+1}) = \mu(z_t) + \sigma(z_t)\left[\rho s_t + \sqrt{1-\rho^2}\,\varepsilon_{t+1}\right], \qquad \varepsilon_{t+1} \sim N(0,1), \tag{17}$$

with $\sigma(z_t) > 0$.

Equation (17) postulates a normal representation intended to approximate the endogenous continuation value. The functions $\mu(z_t)$ and $\sigma(z_t)$ summarize how parental investment affects its mean and its exposure to the common environment. They are maintained objects in this reduced analysis; the recursive fixed point has not been shown to generate this distribution.

Two qualifications are important. Proposition 1 works on a space of bounded value functions, whereas a normal random variable is unbounded. The local model is therefore motivated by the recursion but is not derived from its existence theorem. Establishing that link would require either a larger function space with finite CARA values or a separate approximation result; I do not attempt either here. In addition, s_t changes beliefs about the future and no current object. I hold resources and child costs fixed, $\partial m/\partial s = \partial q/\partial s = 0$, and do not allow s_t to enter μ or σ directly. If bad news also reduced today's resources, fertility could move through the budget even when the belief channel is inactive.

The parameter ρ affects both the mean and residual risk. Conditional on the signal, $\mathbb{E}[V(x_{t+1}, s_{t+1}) \mid s_t, \rho] = \mu(z_t) + \sigma(z_t)\rho s_t$ while $\mathrm{Var}(V(x_{t+1}, s_{t+1}) \mid s_t, \rho) = \sigma(z_t)^2(1-\rho^2)$. Thus a larger ρ gives the news more influence on the conditional mean and leaves less uncertainty unresolved.

Now define the dynastic continuation payoff $Y_{t+1} = \Phi(n_t)V(x_{t+1}, s_{t+1})$. Under (17), Y_{t+1} is conditionally normal. The CARA certainty equivalent is therefore

$$K(n_t, z_t; \rho, s_t) = \Phi(n_t)\left[\mu(z_t) + \sigma(z_t)\rho s_t\right] - \frac{\gamma}{2}\Phi(n_t)^2\sigma(z_t)^2(1-\rho^2). \tag{18}$$

The first term is expected dynastic value and the second is its risk penalty. The penalty contains $\Phi(n_t)^2$ because children share the shock. Adding children increases exposure to one common future rather than diversifying across independent outcomes.

Using (18), the parent's reduced problem is

$$\max_{n,z}\left\{u\Big(m_t - n(q_t+z)\Big) + \beta \min_{\rho\in[\underline{\rho},\overline{\rho}]} K(n, z; \rho, s_t)\right\}. \tag{19}$$

The reduced problem retains the ordering that matters: the parent chooses fertility and investment, evaluates continuation value at the least favorable informativeness, and pays a penalty for unresolved common risk. Fixing $z = \bar{z}$ gives the two-period specialization used for the closed form in Section 3.5; its child payoff should still be read as a stand-in for endogenous continuation value.

3.2 Benchmark with known signal informativeness

Begin with known ρ. Holding the realization s_t fixed, greater informativeness need not move fertility in one direction. It amplifies the content of bad news but also removes residual risk. Proposition 3 separates those effects.

Proposition 3 (Known signal informativeness)**.** *Fix $z_t = z$ and suppose the fertility problem has a unique interior solution $n^*(\rho, s_t)$ satisfying the second-order condition. Assume $\Phi'(n^*) > 0$ and $\sigma(z) > 0$. Under the local normal representation (17),*

$$\operatorname{sgn}\left(\frac{\partial n^*}{\partial \rho}\right) = \operatorname{sgn}\left[s_t + 2\gamma\Phi(n^*)\sigma(z)\rho\right]. \tag{20}$$

Thus higher signal informativeness raises fertility if and only if

$$s_t > -2\gamma\Phi(n^*)\sigma(z)\rho.$$

With endogenous z_t, let H be the Hessian of the objective with respect to (n, z), evaluated at the optimum, and suppose H is negative definite. Then

$$\begin{pmatrix} \dfrac{\partial n^*}{\partial \rho} \\ \dfrac{\partial z^*}{\partial \rho} \end{pmatrix} = -H^{-1} \begin{pmatrix} \beta\Phi'(n^*)\sigma(z^*)\Big[s_t + 2\gamma\Phi(n^*)\sigma(z^*)\rho\Big] \\ \beta\Phi(n^*)\sigma_z(z^*)\Big[s_t + 2\gamma\Phi(n^*)\sigma(z^*)\rho\Big] \end{pmatrix}. \tag{21}$$

Proof. The proof is in Appendix A.4. □

For good news, the mean and risk-resolution effects point in the same direction, so fertility rises with ρ. For bad news they conflict. Risk resolution dominates when $-2\gamma\Phi(n^*)\sigma(z)\rho < s_t < 0$; the adverse mean effect dominates below that range.

When investment is endogenous, the same scalar factor governs both policy responses in (21). At zero, neither fertility nor investment responds locally to a change in known informativeness. Under ambiguity, nature selects ρ by balancing these same forces against the parent.

Figure 1 compares known and ambiguous informativeness. With known ρ, fertility changes smoothly with the signal. Under ambiguity, very bad news is evaluated at $\overline{\rho}$, good news at $\underline{\rho}$, and fertility is flat between them. The right panel displays the adjustment in

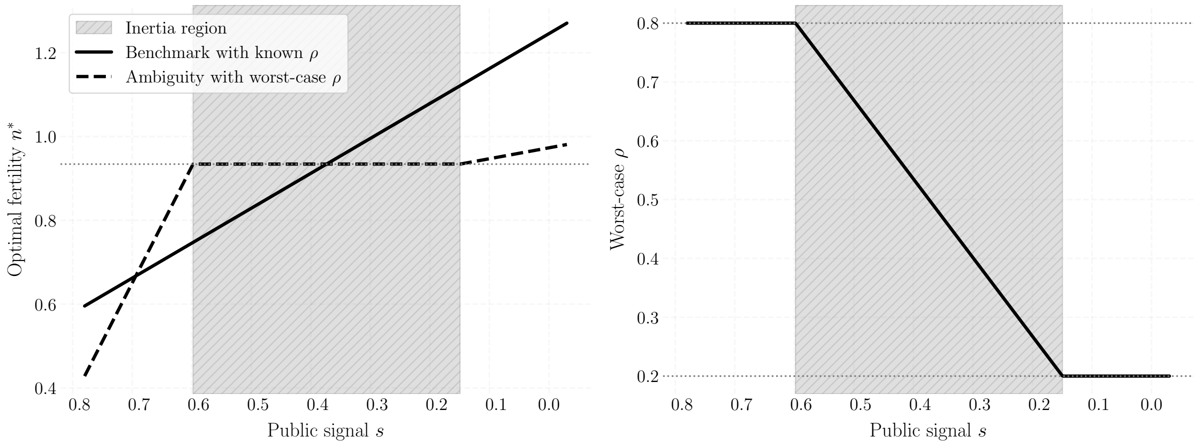

Figure 1: Optimal fertility under known and ambiguous signal informativeness

Notes: The figure compares optimal fertility when signal informativeness is known with optimal fertility when the parent evaluates the signal under the worst admissible informativeness. Parameters are $u(c) = c$, $\Phi(n) = n$, $y = 10$, $q = 1$, $\beta = 0.95$, $\mu = 1.8$, $\sigma = 1$, $\gamma = 0.8$, $\underline{\rho} = 0.2$, and $\overline{\rho} = 0.8$. The known-ρ benchmark uses $\rho = 0.5$. Under these values, the inertia fertility level is $n^0 = 0.934$, and the inertia region is approximately $s \in [-0.598, -0.149]$.

the adverse belief. As news becomes less bad, ρ_t^* falls: the parent exchanges the possibility that bad news is highly relevant for the possibility that more risk remains unresolved. Fertility does not move until the belief reaches an endpoint.

3.3 Ambiguous signal informativeness

I now allow the parent to be unsure how informative the public signal is. Fix a plan (n, z) with $\Phi(n) > 0$. (If $n = 0$, continuation value is zero and K no longer depends on ρ, so the formulas that follow are irrelevant.) For every other plan, the parent uses the least favorable admissible value of ρ and therefore evaluates $\min_{\rho \in [\underline{\rho}, \overline{\rho}]} K(n, z; \rho, s_t)$. Using (18), the derivative with respect to ρ is

$$K_\rho(n, z; \rho, s_t) = \Phi(n)\sigma(z)s_t + \gamma\Phi(n)^2\sigma(z)^2\rho,$$

and

$$K_{\rho\rho}(n, z; \rho, s_t) = \gamma\Phi(n)^2\sigma(z)^2 > 0.$$

Thus K is strictly convex in ρ. Its minimizer is the unconstrained solution, truncated when necessary to remain inside the ambiguity set:

$$\rho^w(n, z, s_t) = \mathrm{Proj}_{[\underline{\rho}, \overline{\rho}]}\left(-\frac{s_t}{\gamma\Phi(n)\sigma(z)}\right). \tag{22}$$

Suppose $0 \le \underline{\rho} < \overline{\rho}$. For good news ($s_t \ge 0$), the unconstrained solution is nonpositive, so the parent settles on $\rho^w = \underline{\rho}$ and gives the signal as little weight as the model permits.

For sufficiently bad news, $s_t \le -\gamma\Phi(n)\sigma(z)\overline{\rho}$, the opposite boundary binds: $\rho^w = \overline{\rho}$. The parent then treats the news as highly informative.

Between these two cases, for

$$-\gamma\Phi(n)\sigma(z)\overline{\rho} < s_t < -\gamma\Phi(n)\sigma(z)\underline{\rho},$$

the worst case is interior so that

$$\rho^w(n, z, s_t) = -\frac{s_t}{\gamma\Phi(n)\sigma(z)}.$$

This middle range is where the mechanism does its work. Here the least-favorable interpretation varies with both the news and the plan being considered. Planning a larger family, or choosing an investment that raises $\sigma(z)$, changes which value of ρ is most damaging. The result is asymmetric: the parent discounts favorable news and gives adverse news more credence. Fertility inertia arises in the interior range, where that pessimistic interpretation can move freely.

Proposition 4 (Fertility inertia)**.** *Let (n^0, z^0) be the unique interior maximizer of the signal-free problem*

$$\max_{n,z} \left\{ u\Big(m_t - n(q_t + z)\Big) + \beta \left[\Phi(n)\mu(z) - \frac{\gamma}{2}\Phi(n)^2\sigma(z)^2\right] \right\}. \tag{23}$$

Suppose the objective in (23) is strictly concave in a neighborhood of (n^0, z^0), with $\Phi(n^0) > 0$ and $\sigma(z^0) > 0$. If the signal satisfies

$$-\gamma\Phi(n^0)\sigma(z^0)\overline{\rho} < s_t < -\gamma\Phi(n^0)\sigma(z^0)\underline{\rho}, \tag{24}$$

then the worst-case signal informativeness is interior at (n^0, z^0), and the optimal policies in (19) satisfy, locally,

$$n^*(s_t) = n^0, \qquad z^*(s_t) = z^0, \qquad c^*(s_t) = m_t - n^0(q_t + z^0). \tag{25}$$

Thus, throughout the inertia region, optimal fertility, child investment, and consumption are locally independent of the public signal.

Proof. The proof is available in Appendix A.5. □

The proposition is deliberately local. It identifies (n^0, z^0) as the only nearby optimum among plans with an interior worst-case belief. The same pair is globally optimal if every candidate plan has an interior worst-case belief and (23) is strictly concave over the full choice set; these conditions exclude rival optima at $n = 0$ or at a boundary value of ρ. The endpoints in (24) require a little more care. At either endpoint, the worst-case belief just touches the edge of the ambiguity set. A nearby deviation then receives a small convex gain, so the strengthened second-order condition in the appendix is needed to ensure that local concavity still dominates. In the one-dimensional example in Section 3.5, this condition is automatic.

3.4 The welfare cost of ambiguous news

The source of the flat policy response becomes transparent once the interior minimizer in (22) is substituted into (18). The resulting continuation value is

$$\min_{\rho \in [\underline{\rho}, \overline{\rho}]} K(n, z; \rho, s_t) = \Phi(n)\mu(z) - \frac{\gamma}{2}\Phi(n)^2\sigma(z)^2 - \frac{s_t^2}{2\gamma}. \tag{26}$$

The first two terms are what the parent would obtain from an uninformative signal: expected dynastic value less the risk penalty associated with the unresolved common component. The last term is the welfare loss from ambiguous news. Crucially, it does not vary across plans as long as their worst-case beliefs remain interior. The news makes the parent worse off, but it does not alter the marginal return to another child or to child investment. The first-order conditions inside the inertia region are consequently identical to those of the signal-free problem. Household choices appear unresponsive even though the signal matters for welfare.

Corollary 1 (Width of the inertia region)**.** *For a fixed signal-free optimum (n^0, z^0), the length of the fertility-inertia interval is*

$$\gamma\, \Phi(n^0)\, \sigma(z^0)\, (\overline{\rho} - \underline{\rho}). \tag{27}$$

Thus the direct width of the inertia region is larger when dynastic risk sensitivity γ, fertility exposure $\Phi(n^0)$, continuation-value exposure $\sigma(z^0)$, or ambiguity about signal informativeness $\overline{\rho} - \underline{\rho}$ is larger.

Each factor in (27) has a simple role. Greater risk sensitivity makes unresolved risk more costly; higher $\Phi(n^0)$ or $\sigma(z^0)$ puts more dynastic value at stake; and a wider ambiguity set allows the pessimistic belief to adjust over a larger range. Holding the signal-free plan fixed, all three enlarge the interval on which the parent keeps that plan. If a parameter change also shifts n^0 or z^0, those indirect effects must of course be included in the total comparative static. The flat portion lies on the bad-news side: good news is read using $\underline{\rho}$, the worst news using $\overline{\rho}$, and choices are locally constant between those cases.

3.5 Closed-form special case

To see the policy rule without the surrounding notation, consider the simplest two-period version of the model. Set $u(c) = c$ and $\Phi(n) = n$, and hold child investment fixed at $z = 0$. Then

$$\mu \equiv \mu(0), \qquad \sigma \equiv \sigma(0), \qquad q_t = q, \qquad y \equiv m_t.$$

Here q denotes the resource cost of a child and y the parent's current resources.

In the absence of a signal, fertility solves

$$\max_{0 \le n \le y/q} \left\{ y - qn + \beta \left[n\mu - \frac{\gamma}{2} n^2 \sigma^2 \right] \right\}.$$

The objective is strictly concave because the quadratic term in n is negative

$$\frac{\partial^2}{\partial n^2} \left\{ y - qn + \beta \left[n\mu - \frac{\gamma}{2} n^2 \sigma^2 \right] \right\} = -\beta\gamma\sigma^2 < 0.$$

The interior first-order condition therefore gives

$$n^0 = \frac{\mu - q/\beta}{\gamma\sigma^2}. \tag{28}$$

This solution is interior when

$$\mu > q/\beta \qquad \text{and} \qquad qn^0 < y.$$

The first inequality requires the mean dynastic payoff from an additional child to exceed its discount-adjusted cost. The second simply ensures that the implied family size is affordable.

Define

$$\lambda \equiv \frac{\mu - q/\beta}{\sigma}.$$

I use λ as a convenient risk-adjusted measure of the attractiveness of parenthood: it is the excess mean payoff per unit of common risk. By (28),

$$\gamma\sigma n^0 = \lambda.$$

Therefore the inertia interval becomes

$$s_t \in [-\lambda\overline{\rho},\, -\lambda\underline{\rho}], \qquad \text{with width} \qquad \lambda(\overline{\rho} - \underline{\rho}). \tag{29}$$

Equation (29) separates two sources of inertia. The first, λ, captures the risk-adjusted appeal of parenthood and hence the size of the parent's exposure. The second, $\overline{\rho} - \underline{\rho}$, measures how much latitude there is to reinterpret the signal. The flat segment disappears when ρ is known and grows linearly as the ambiguity interval widens around a fixed midpoint. Thus the relevant friction is uncertainty about informativeness, rather than merely a signal that happens to be weak.

The linear specification also yields closed forms outside the inertia region, provided that $0 < n^*(s_t) < y/q$. Extremely bad signals eventually push fertility to zero, whereas sufficiently good signals make the resource constraint bind. At either corner, fertility is flat for the mechanical reason that a constraint, rather than the ambiguity mechanism, determines the choice. When the worst-case belief is fixed at a boundary $\hat{\rho} \in \{\underline{\rho}, \overline{\rho}\}$, the interior solution is

$$n^*(s_t; \hat{\rho}) = \frac{\mu + \sigma\hat{\rho}s_t - q/\beta}{\gamma\sigma^2(1 - \hat{\rho}^2)}.$$

Corollary 2 (Closed-form policy function)**.** *Assume the specification of this subsection with* $0 \leq \underline{\rho} < \overline{\rho} < 1$*, and suppose the solution remains interior,* $0 < n^*(s_t) < y/q$*. Then the ambiguous-signal policy is*

$$n^*(s_t) = \begin{cases} \dfrac{\mu + \sigma\overline{\rho}s_t - q/\beta}{\gamma\sigma^2(1 - \overline{\rho}^2)}, & s_t \leq -\lambda\overline{\rho}, \\[2em] n^0, & -\lambda\overline{\rho} \leq s_t \leq -\lambda\underline{\rho}, \\[2em] \dfrac{\mu + \sigma\underline{\rho}s_t - q/\beta}{\gamma\sigma^2(1 - \underline{\rho}^2)}, & s_t \geq -\lambda\underline{\rho}. \end{cases} \tag{30}$$

The three pieces paste together continuously. The policy is flat on the inertia region, and outside the region its slope is

$$\frac{\partial n^*(s_t; \hat{\rho})}{\partial s_t} = \frac{\hat{\rho}}{\gamma\sigma(1 - \hat{\rho}^2)},$$

with $\hat{\rho} = \overline{\rho}$ *on the bad-news side and* $\hat{\rho} = \underline{\rho}$ *on the good-news side.*

The three pieces meet continuously but with kinks. At $s_t = -\lambda\overline{\rho}$, the bad-news branch reaches n^0; at $s_t = -\lambda\underline{\rho}$, the good-news branch does the same. Because $\rho/(1 - \rho^2)$ rises on $[0, 1)$, fertility reacts more sharply on the bad-news branch, governed by $\overline{\rho}$, than on the good-news branch, governed by $\underline{\rho}$. This is the model's asymmetry in its clearest form: parents act strongly on sufficiently adverse news while discounting favorable news.

3.6 Generosity and legibility of family policy

Up to this point, policy has appeared only as part of the public signal. I now introduce a family benefit directly. The central result is that a household responds to what it can safely expect to receive, which need not bear much relation to the program's headline cost.

Consider a statutory subsidy $\tau \geq 0$ for each child. Its announced value is not necessarily the amount a household expects to collect over the years in which it will raise that child. Eligibility may be means-tested, childcare places may be rationed, applications may be burdensome, or a later government may change the program. Let α denote the fraction of the statutory benefit that ultimately reaches the household. Rather than assigning a prior to α, the household can only place it in the interval

$$\alpha \in [\underline{\alpha}, \overline{\alpha}] \subseteq [0, 1], \tag{31}$$

just as it does with signal informativeness. In this notation, a program is more legible in the specific sense of having a higher minimum credible pass-through $\underline{\alpha}$. Narrowing the interval by lowering $\overline{\alpha}$ alone has no effect on behavior in this specification. A clearer information environment likewise has a narrower $[\underline{\rho}, \overline{\rho}]$. Because the net cost of a child is $q - \alpha\tau$, the closed-form parent in Section 3.5 solves

$$\max_{0 \leq n \leq y/(q-\underline{\alpha}\tau)} \min_{\substack{\rho \in [\underline{\rho}, \overline{\rho}] \\ \alpha \in [\underline{\alpha}, \overline{\alpha}]}} \left\{ y - (q - \alpha\tau)n + \beta \left[n(\mu + \sigma\rho s_t) - \frac{\gamma}{2} n^2 \sigma^2 (1 - \rho^2) \right] \right\}. \tag{32}$$

For any $n > 0$ and $\tau > 0$, the objective rises with α. The least-favorable evaluation therefore sets $\alpha = \underline{\alpha}$, regardless of the family plan or the value chosen for ρ. The two minimizations separate, and the earlier results carry over after replacing q with

$$q_e(\tau) \equiv q - \underline{\alpha}\tau.$$

Proposition 5 (Generosity and legibility)**.** *Consider problem (32) with* $0 \leq \underline{\alpha} \leq \overline{\alpha} \leq 1$ *and* $q - \underline{\alpha}\tau > 0$*, and suppose the signal-free solution is interior. Then the following hold.*

(i) The worst-case pass-through is $\alpha^* = \underline{\alpha}$ *for every plan with* $n > 0$*, and all results of Sections 3.1–3.5 hold with* q *replaced by* $q_e(\tau)$*. In particular, inside the inertia region fertility equals*

$$n^0(\tau) = \frac{\mu - \big(q - \underline{\alpha}\tau\big)/\beta}{\gamma\sigma^2}. \tag{33}$$

(ii) The fertility response to the subsidy is

$$\frac{\partial n^0}{\partial \tau} = \frac{\underline{\alpha}}{\beta\gamma\sigma^2}. \tag{34}$$

The response is scaled by the worst-case pass-through $\underline{\alpha}$*, not by the statutory size* τ*.*

(iii) (Legibility equivalence.) An ambiguous program $(\tau, [\underline{\alpha}, \overline{\alpha}])$ *has the same effect on every household decision as a fully legible program* ($\alpha \equiv 1$*, known) of size* $\underline{\alpha}\tau$*. Equivalently, for* $\underline{\alpha} > 0$*, replicating the effect of a legible program of size* τ^* *requires scaling the ambiguous program up by the factor* $1/\underline{\alpha}$*. For* $\underline{\alpha} = 0$*, no finite statutory subsidy replicates a positive legible one.*

(iv) The width of the inertia region, $\lambda(\tau)(\overline{\rho} - \underline{\rho})$ *with* $\lambda(\tau) = \big[\mu - (q - \underline{\alpha}\tau)/\beta\big]/\sigma$*, is nondecreasing in* τ*,*

$$\frac{\partial}{\partial \tau}\Big[\lambda(\tau)(\overline{\rho} - \underline{\rho})\Big] = \frac{\underline{\alpha}}{\beta\sigma}(\overline{\rho} - \underline{\rho}) \geq 0,$$

with strict inequality if and only if $\underline{\alpha} > 0$ *and* $\overline{\rho} > \underline{\rho}$*.*

Proof. The proof is available in Appendix A.6. □

Part (ii) includes a stark limiting case. As $\underline{\alpha} \to 0$, the fertility response vanishes for any statutory benefit, even though government spending need not. If the realized pass-through is some $\alpha_0 \in [\underline{\alpha}, \overline{\alpha}]$ with $\alpha_0 > 0$, the government pays $g = \alpha_0 \tau$ per child, and the fertility response per unit of realized per-child transfer is

$$\frac{\partial n^0}{\partial g} = \frac{\underline{\alpha}}{\alpha_0} \cdot \frac{1}{\beta\gamma\sigma^2} \;\leq\; \frac{1}{\beta\gamma\sigma^2},$$

which is the response to a fully legible transfer multiplied by $\underline{\alpha}/\alpha_0 \leq 1$. Put plainly, the household plans against the lower bound, while the treasury pays according to the realized program.

This wedge identifies a channel through which a program can generate substantial outlays and a small fertility response. Its size depends on the ratio of the decision-relevant lower bound, $\underline{\alpha}$, to realized pass-through, α_0. Administrative or political uncertainty matters here insofar as it lowers the amount the household can count on receiving. The empirical literature discussed in the introduction motivates this possibility but does not identify either bound or quantify this channel.

Part (iv) points to a second, less obvious effect. A subsidy raises the value at stake in the parent's dynastic position. The parent then becomes cautious over a wider range of news, so the inertia interval expands. Spending more through a program that families do not trust can therefore make the response to the next announcement even weaker.

The announcement itself is also a public signal. A credible new program is good news, but Corollary 2 says that the parent reads good news using $\underline{\rho}$. Its effect on fertility therefore has the muted slope $\underline{\rho}/[\gamma\sigma(1 - \underline{\rho}^2)]$. In practice, governments often announce family packages after demographic or economic conditions have deteriorated. The package

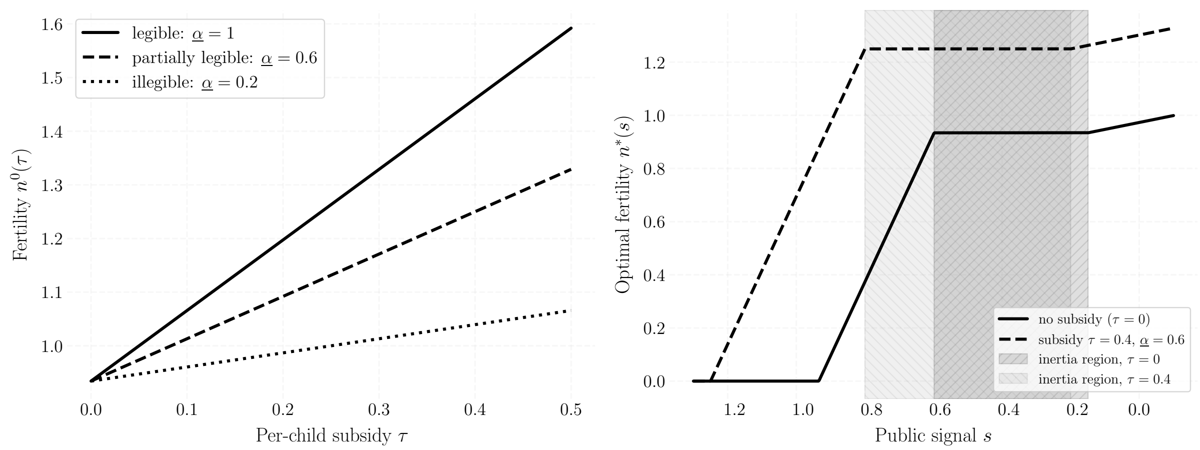

Figure 2: Generosity versus legibility of family policy

Notes: The left panel plots inertia-level fertility $n^0(\tau)$ in (33) against the statutory per-child subsidy τ for three levels of worst-case pass-through $\underline{\alpha}$. The slope is $\underline{\alpha}/(\beta\gamma\sigma^2)$, so the same statutory generosity produces responses ranked by legibility. The right panel plots the full policy function $n^*(s)$ without a subsidy and with a subsidy $\tau = 0.4$ under $\underline{\alpha} = 0.6$. The subsidy raises the inertia level but also shifts and widens the inertia region (Proposition 5(iv)). Parameter values are $u(c) = c$, $\Phi(n) = n$, $y = 10$, $q = 1$, $\beta = 0.95$, $\mu = 1.8$, $\sigma = 1$, $\gamma = 0.8$, $\underline{\rho} = 0.2$, and $\overline{\rho} = 0.8$.

and the circumstances that prompted it may together fall inside $[-\lambda\overline{\rho}, -\lambda\underline{\rho}]$. Fertility then does not move at all, even though households have paid attention. An event-study coefficient near zero need not imply that the announcement went unnoticed.

The natural policy response is to make the program easier to read and rely on. Such a reform raises $\underline{\alpha}$ toward $\overline{\alpha}$, because a household can verify what it will receive, and narrows $[\underline{\rho}, \overline{\rho}]$, because public conditions become easier to map into the household's own future. These changes affect the inertia width $W = \lambda(\tau, \underline{\alpha})(\overline{\rho} - \underline{\rho})$ through different channels. Raising $\underline{\alpha}$ improves the pass-through of each dollar already on offer. When $\tau > 0$, however, it also raises the dynastic value at stake and, by the force in Proposition 5(iv), tends to widen W: $\partial W/\partial\underline{\alpha} = \tau(\overline{\rho} - \underline{\rho})/(\beta\sigma) \geq 0$. By contrast, narrowing the informativeness interval unambiguously reduces the width: $\partial W/\partial(\overline{\rho} - \underline{\rho}) = \lambda > 0$. The total effect is

$$dW = \frac{\tau(\overline{\rho} - \underline{\rho})}{\beta\sigma}\, d\underline{\alpha} + \lambda\, d(\overline{\rho} - \underline{\rho}).$$

The inertia region shrinks only when the reduction in $\overline{\rho} - \underline{\rho}$ outweighs the exposure effect of the higher pass-through. Thus legibility always increases the return to a given subsidy, but it restores responsiveness to news only if its informational component is large enough. Figure 2 shows the pass-through channel. Guaranteed rather than rationed childcare, statutory eligibility, enforceable return-to-work rights, and benefit schedules fixed for a credible horizon all operate on the bounds themselves, not just on τ. They make the promised payoff easier to verify and more likely to arrive. The model therefore predicts that a specific, dependable benefit can move births at a lower fiscal cost than a larger but poorly defined promise.

3.7 How special is exact flatness?

Exact flatness in the benchmark relies on maxmin preferences and a continuous fertility choice. I consider two departures—smooth ambiguity and integer fertility—to see which parts of the mechanism remain. These are targeted robustness exercises, not a claim that ambiguity is necessary for every form of inaction. Appendix B gives the details.

Under smooth ambiguity (Klibanoff et al., 2005), the parent has a prior π over $[\underline{\rho}, \overline{\rho}]$ and an ambiguity-aversion parameter η. As η grows, the value criterion approaches maxmin, and the endogenous tilt over ρ points in the same direction as the worst-case selection: worse news places more weight on less favorable interpretations (Appendix B.1). Exact flatness appears only in the limit $\eta = \infty$. With finite η, the policy is smooth, and its slope contains a covariance term whose sign is not fixed. One therefore cannot assert monotonicity or a universal S-shape without more structure. Convergence of the policy itself, as opposed to its value, also requires the additional conditions stated in the appendix.

Now suppose children come only in whole numbers (Proposition B.2.1). The parent compares neighboring plans, such as n and $n + 1$. If the worst-case belief is interior for both, their values contain the same penalty, $-s_t^2/(2\gamma)$, which cancels from the comparison. Bad news lowers welfare without changing which family size is preferred. Except at an exact tie, the parent cannot switch plans anywhere in this common interior region. Switching occurs only after a worst-case belief reaches a boundary; $\overline{\rho}$ then governs the bad-news side and $\underline{\rho}$ the good-news side.

This distinction suggests a possible empirical test: assessments of future conditions could worsen while the chosen positive family size remains unchanged. Such a test would require a measurement link between reported assessments and model value; fertility intentions are not themselves a measure of welfare. Integer choice by itself creates intervals of inaction, so I do not claim that ambiguity always makes those intervals wider. Relative to a known-ρ model, their width can increase or decrease depending on the comparison value of ρ. The robust implication is instead the placement of the switching points and the separation between welfare and behavior while both plans have interior worst-case beliefs.

4 Conclusion

Ambiguity about the relevance of public information can separate welfare from fertility choice. In the local normal model, the least-favorable informativeness parameter adjusts to moderately adverse news so that the resulting loss in value is independent of the family plan. Fertility, consumption, and investment per child then remain locally constant even as the parent's evaluation deteriorates. The closed-form specialization places this interval between a steeper response to sufficiently bad news and a shallower response to good news. Smooth ambiguity and positive integer family sizes preserve specified aspects of the mechanism under the conditions established in the appendices.

Insensitivity to information can coexist with sensitivity to material incentives. The policy application distinguishes an announced transfer from its minimum credible household-level delivery. An ambiguous subsidy has the same effect as a certain subsidy scaled by its lowest admissible pass-through, while realized spending can reflect a higher delivery rate. Increasing the effective subsidy also raises common-risk exposure and can widen the inertia interval. Improving benefit delivery and narrowing ambiguity about public information therefore have different implications for responsiveness to subsequent news.

The empirical findings discussed in the introduction motivate these questions; the paper's contribution is their theoretical analysis. Further work could establish conditions linking the local normal representation to the recursive value function and study learning about information quality. Empirical tests would need to distinguish news about future conditions from changes in current resources and to measure uncertainty about benefit delivery separately from statutory generosity.

Acknowledgements

I am deeply grateful to my advisor, Juan Carlos Córdoba, for invaluable advice, guidance, and encouragement. I am grateful to my dissertation committee, Rajesh Singh, Quinn Weninger, Sergio Lence, and Jian Li, for their continuous support. I am also grateful to audiences at the Midwest Economic Association (2023), the Midwest Macro Meetings at Purdue (2024), the 19th Economics Graduate Student Conference at Washington University in St. Louis (2024), and the 94th Southern Economic Association meetings (2024), and to seminar participants at the Iowa State reading group, for useful comments and suggestions.

Declarations

Funding The author declares that no funds, grants, or other support were received during the preparation of this manuscript.

Competing interests The author has no relevant financial or non-financial interests to disclose.

Ethics approval and consent to participate Not applicable. This article does not contain studies with human participants or animals performed by the author.

Consent for publication Not applicable.

Data availability This theoretical article uses no proprietary data. All figures are generated by the model, and replication code is available from the author upon request.

Materials availability Not applicable.

Code availability Figure-generation and replication code is available from the author upon request.

Author contribution HongSeok Kim is the sole author and completed all aspects of the work.

Appendices

A Proofs

A.1 Derivation of CARA

Given the future net payoff of a marginal child X_{t+1}, the conditional mean

$$\mathbb{E}[X_{t+1}|s,\rho] = \mu + \sigma\rho s$$

and

$$\mathrm{Var}(X_{t+1} \mid s,\rho) = \sigma^2(1-\rho^2).$$

The total payoff of the parent having n children is

$$Z_{t+1} = \Phi(n)X_{t+1}$$

where $\Phi(n)$ is the effective dynastic discounting factor. Conditional on signal s and parameter ρ,

$$Z_{t+1}|s,\rho \sim N(m,v)$$

where

$$m = \Phi(n)[\mu + \sigma\rho s]$$

and

$$v = \Phi(n)^2\sigma^2(1-\rho^2).$$

Given that the utility function is CARA, the parent evaluates the risky future dynastic payoff Z,

$$U(Z) = -\exp(-\gamma Z)$$

where $\gamma > 0$ is absolute risk aversion. Let

$$-\exp(-\gamma CE) = \mathbb{E}(-\exp(-\gamma Z))$$

and using $Z \sim N(m,v)$, we get

$$\mathbb{E}\left[\exp(-\gamma Z)\right] = \exp\left(-\gamma m + \frac{1}{2}\gamma^2 v\right).$$

Therefore,

$$\exp(-\gamma CE) = \exp\left(-\gamma m + \frac{1}{2}\gamma^2 v\right)$$

which gives

$$CE = m - \frac{\gamma}{2}v = \mathbb{E}[Z] - \frac{\gamma}{2}\operatorname{Var}(Z).$$

Substituting the conditional mean and variance of random variable Z yields the equation

$$CE(n; s, \rho) = \Phi(n)[\mu + \sigma\rho s] - \frac{\gamma}{2}\Phi(n)^2\sigma^2(1-\rho^2).$$

A.2 Proof of Proposition 1

Proof. Let $C_b(\mathcal{X} \times \mathcal{S})$ be the space of bounded continuous functions on the state space, with the sup norm $\|\cdot\|_\infty$, a complete metric space. Write a feasible action as $a = (c, n, z) \in \Gamma(x, s)$ and $V' = V(x', s')$. We show T maps C_b into itself and is a contraction with modulus $\beta\overline{\Phi}$, then characterize the fixed point by measurable selections.

We use two properties of $\mathcal{C}_{\gamma,\rho,z}$, both immediate from $\mathcal{C}_{\gamma,\rho,z}[Y] = -\gamma^{-1}\log\mathbb{E}_{\rho,z}[e^{-\gamma Y}]$ (and from the conditional expectation when $\gamma = 0$). The properties are monotonicity, $Y_1 \leq Y_2 \Rightarrow \mathcal{C}_{\gamma,\rho,z}[Y_1] \leq \mathcal{C}_{\gamma,\rho,z}[Y_2]$, and translation invariance, $\mathcal{C}_{\gamma,\rho,z}[Y + k] = \mathcal{C}_{\gamma,\rho,z}[Y] + k$ for constant k. Together they give the sup-norm bound. If $\|Y_1 - Y_2\|_\infty \leq \eta$ then

$$\left|\mathcal{C}_{\gamma,\rho,z}[Y_1] - \mathcal{C}_{\gamma,\rho,z}[Y_2]\right| \leq \eta. \tag{$*$}$$

***T* maps C_b into C_b.** Fix $V \in C_b$. By Assumption 1, Γ is nonempty, compact-valued, and continuous, $[\underline{\rho}, \overline{\rho}]$ is compact, u is continuous and bounded on the graph of Γ ($|u| \leq \overline{u}$), and $0 \leq \Phi(n) \leq \overline{\Phi}$. Then $|\Phi(n)V'| \leq \overline{\Phi}\|V\|_\infty$, and applying ($*$) against the constant bounds gives $|\mathcal{C}_{\gamma,\rho,z}[\Phi(n)V']| \leq \overline{\Phi}\|V\|_\infty$ uniformly in ρ, hence $|(TV)(x,s)| \leq \overline{u} + \beta\overline{\Phi}\|V\|_\infty$, so TV is bounded. For continuity, set $H_V(x,s,a,\rho) = u(x,s,a) + \beta\mathcal{C}_{\gamma,\rho,z}[\Phi(n)V']$. Under Assumption 1 the transition law is continuous and the integrand $e^{-\gamma\Phi(n)V'}$ is bounded and continuous, so by bounded convergence H_V is continuous in (x,s,a,ρ). Since $[\underline{\rho}, \overline{\rho}]$ is compact and fixed in (x,s,a), Berge's theorem makes $G_V(x,s,a) := \min_\rho H_V(x,s,a,\rho)$ continuous with nonempty compact-valued argmin. Applying Berge's theorem again with the compact-valued continuous Γ makes $(TV)(x,s) = \max_{a\in\Gamma(x,s)} G_V(x,s,a)$ continuous.

***T* is a contraction.** Let $V_1, V_2 \in C_b$ and $\Delta := \|V_1 - V_2\|_\infty$. For fixed feasible (x,s,a,ρ), $|V_1' - V_2'| \leq \Delta$ and $\Phi(n) \geq 0$ give $\|\Phi(n)V_1' - \Phi(n)V_2'\|_\infty \leq \overline{\Phi}\Delta$, so by ($*$),

$$\left|\mathcal{C}_{\gamma,\rho,z}[\Phi(n)V_1'] - \mathcal{C}_{\gamma,\rho,z}[\Phi(n)V_2']\right| \leq \overline{\Phi}\Delta$$

uniformly in ρ. The minimization over ρ and the maximization over a are each nonexpansive in the sup norm. If $\sup_\rho |f_\rho - g_\rho| \leq \eta$ then $|\min_\rho f_\rho - \min_\rho g_\rho| \leq \eta$, and likewise for max (from $f \leq g + \eta$, take the relevant extremum of both sides, then swap roles). The payoff $u(x,s,a)$ cancels in the difference. Passing the bound through the inner min, multiplying by β, and passing through the outer max gives, for every (x, s),

$$|(TV_1)(x, s) - (TV_2)(x, s)| \leq \beta\overline{\Phi}\Delta,$$

hence $\|TV_1 - TV_2\|_\infty \leq \beta\overline{\Phi}\|V_1 - V_2\|_\infty$. Since $\beta\overline{\Phi} < 1$ by Assumption 1, T is a contraction on C_b, and Banach's fixed-point theorem gives a unique $V \in C_b$ with $V = TV$.

Characterization. At this fixed point define

$$\begin{aligned}\mathcal{R}_V(x, s, a) &:= \arg\min_{\rho\in[\underline{\rho},\overline{\rho}]}\mathcal{C}_{\gamma,\rho,z}\Big[\Phi(n)V'\Big],\\ \mathcal{P}_V(x, s) &:= \arg\max_{a\in\Gamma(x,s)}\Big\{u(x, s, a) + \beta\min_{\rho}\mathcal{C}_{\gamma,\rho,z}\Big[\Phi(n)V'\Big]\Big\}.\end{aligned}$$

By the continuity and compactness above, Berge's theorem makes both correspondences nonempty and compact-valued, each admitting a measurable selection. For any selections $a^*(x, s) = (c^*, n^*, z^*)(x, s) \in \mathcal{P}_V(x, s)$ and $\rho^*(x, s) \in \mathcal{R}_V(x, s, a^*(x, s))$,

$$V(x, s) = u\Big(x, s, a^*(x, s)\Big) + \beta\,\mathcal{C}_{\gamma,\rho^*(x,s),z^*(x,s)}\Big[\Phi(n^*(x, s))V'\Big] \qquad \text{for every } (x, s).$$

The worst-case belief $\rho^*(x, s)$ is the inner minimizer evaluated at the chosen policy $a^*(x, s)$, which is the belief endogeneity used in the inertia analysis. □

A.3 Proof of Proposition 2

Proof. Fix the current state (x_t, s_t) and suppress it where no confusion can arise. Let

$$J(n, z; x_t, s_t) := \min_{\rho\in[\underline{\rho},\overline{\rho}]} K_\rho(n, z; x_t, s_t).$$

The parent's problem can be written as

$$\max_{c,n,z}\{u(c) + \beta J(n, z; x_t, s_t)\}$$

subject to

$$c + n(q_t + z) \leq m_t.$$

Since $u'(c) > 0$, the budget constraint binds at an optimum. Hence

$$c_t^* = m_t - n_t^*(q_t + z_t^*),$$

which gives (12).

Let λ_t denote the multiplier on the budget constraint. For an interior solution, the first-order conditions are

$$u'(c_t^*) = \lambda_t,$$

$$\beta J_n(n_t^*, z_t^*; x_t, s_t) = \lambda_t(q_t + z_t^*),$$

and

$$\beta J_z(n_t^*, z_t^*; x_t, s_t) = \lambda_t n_t^*.$$

Substituting $\lambda_t = u'(c_t^*)$ gives

$$u'(c_t^*)(q_t + z_t^*) = \beta J_n(n_t^*, z_t^*; x_t, s_t)$$

and

$$n_t^* u'(c_t^*) = \beta J_z(n_t^*, z_t^*; x_t, s_t).$$

It remains to compute the derivatives of J. For fixed ρ, use the differentiable shock representation

$$x_{t+1} = G^x(x_t, s_t, z_t, \varepsilon_{t+1}; \rho), \qquad s_{t+1} = G^s(x_t, s_t, z_t, \varepsilon_{t+1}; \rho),$$

where the conditional law of ε_{t+1} depends on ρ but not on n or z (Assumption 1(v)). This common-shock structure is what keeps the derivative below free of a likelihood-score term. If instead the law of the shock varied directly with z, with density $p_{\rho,z}$, the investment derivative would contain the additional term $-\gamma^{-1}\mathbb{E}_Q[\partial_z \log p_{\rho,z}]$. Define

$$Y_{t+1}(n, z, \rho) := \Phi(n) V\left(G^x(x_t, s_t, z, \varepsilon_{t+1}; \rho), G^s(x_t, s_t, z, \varepsilon_{t+1}; \rho)\right).$$

For $\gamma > 0$,

$$K_\rho(n, z; x_t, s_t) = -\frac{1}{\gamma} \log \mathbb{E}_\rho\left[\exp\{-\gamma Y_{t+1}(n, z, \rho)\} \mid x_t, s_t\right],$$

where the expectation is over the shock ε_{t+1}.

By the maintained differentiability and domination assumptions, differentiation under the expectation is valid. For any control $a \in \{n, z\}$,

$$\frac{\partial K_\rho}{\partial a} = -\frac{1}{\gamma} \frac{\mathbb{E}_\rho\left[-\gamma \exp\{-\gamma Y_{t+1}\} \frac{\partial Y_{t+1}}{\partial a} \mid x_t, s_t\right]}{\mathbb{E}_\rho\left[\exp\{-\gamma Y_{t+1}\} \mid x_t, s_t\right]} = \frac{\mathbb{E}_\rho\left[\exp\{-\gamma Y_{t+1}\} \frac{\partial Y_{t+1}}{\partial a} \mid x_t, s_t\right]}{\mathbb{E}_\rho\left[\exp\{-\gamma Y_{t+1}\} \mid x_t, s_t\right]}.$$

Define the distorted measure $Q_{\rho,n,z}$ by

$$\frac{\mathrm{d}Q_{\rho,n,z}}{\mathrm{d}P_{\rho,z}} = \frac{\exp\{-\gamma Y_{t+1}(n, z, \rho)\}}{\mathbb{E}_{\rho,z}\left[\exp\{-\gamma Y_{t+1}(n, z, \rho)\} \mid x_t, s_t\right]}.$$

Equivalently, under the shock representation, this is the exponential tilt of the shock law.

Therefore

$$\frac{\partial K_\rho}{\partial a} = \mathbb{E}_{Q_{\rho,n,z}}\left[\frac{\partial Y_{t+1}}{\partial a} \mid x_t, s_t\right].$$

Since the inner minimizer ρ_t^* is unique, Danskin's envelope theorem (Danskin, 1966; Milgrom and Segal, 2002) implies

$$J_a(n_t^*, z_t^*; x_t, s_t) = \left.\frac{\partial K_\rho(n_t^*, z_t^*; x_t, s_t)}{\partial a}\right|_{\rho=\rho_t^*}, \qquad a \in \{n, z\}.$$

Thus there is no additional term involving the derivative of ρ_t^* with respect to n or z.

Now compute the two derivatives of Y_{t+1}. Since the transition does not depend directly on n,

$$\frac{\partial Y_{t+1}}{\partial n} = \Phi'(n_t^*)V(x_{t+1}, s_{t+1}).$$

Therefore

$$J_n(n_t^*, z_t^*; x_t, s_t) = \Phi'(n_t^*)\mathbb{E}_{Q_t^*}\left[V(x_{t+1}, s_{t+1}) \mid x_t, s_t\right],$$

where $Q_t^* = Q_{\rho_t^*, n_t^*, z_t^*}$. Substituting into the first-order condition for n gives

$$u'(c_t^*)(q_t + z_t^*) = \beta\Phi'(n_t^*)\mathbb{E}_{Q_t^*}\left[V(x_{t+1}, s_{t+1}) \mid x_t, s_t\right],$$

which is (13).

For z, the chain rule gives

$$\begin{aligned}\frac{\partial Y_{t+1}}{\partial z} &= \Phi(n_t^*)\left[\nabla_x V(x_{t+1}, s_{t+1}) \cdot G_z^x(x_t, s_t, z_t^*, \varepsilon_{t+1}; \rho_t^*)\right.\\ &\qquad\left. + V_s(x_{t+1}, s_{t+1})G_z^s(x_t, s_t, z_t^*, \varepsilon_{t+1}; \rho_t^*)\right]\\ &= \Phi(n_t^*)D_z V_{t+1}.\end{aligned}$$

Hence

$$J_z(n_t^*, z_t^*; x_t, s_t) = \Phi(n_t^*)\mathbb{E}_{Q_t^*}\left[D_z V_{t+1} \mid x_t, s_t\right].$$

Substituting into the first-order condition for z gives

$$n_t^* u'(c_t^*) = \beta\Phi(n_t^*)\mathbb{E}_{Q_t^*}\left[D_z V_{t+1} \mid x_t, s_t\right],$$

which is (14).

Finally, by definition of the inner problem,

$$\rho_t^* \in \arg\min_{\rho \in [\underline{\rho}, \overline{\rho}]} K_\rho(n_t^*, z_t^*; x_t, s_t).$$

If the minimizer is interior and K_ρ is differentiable in ρ, the usual first-order condition gives $\partial_\rho K_\rho|_{\rho=\rho_t^*} = 0$. If $\rho_t^* = \underline{\rho}$ the right derivative must satisfy $\partial_\rho^+ K_\rho|_{\rho=\underline{\rho}} \geq 0$, and if $\rho_t^* = \overline{\rho}$ the left derivative must satisfy $\partial_\rho^- K_\rho|_{\rho=\overline{\rho}} \leq 0$, which is (15). This proves the result. □

A.4 Proof of Proposition 3

Proof. Fix $z_t = z$ and write

$$c(n) = m_t - n(q_t + z).$$

The fertility first-order condition is

$$G(n, \rho; s_t) \equiv -(q_t + z)u'(c(n)) + \beta\Phi'(n)\left[\mu(z) + \sigma(z)\rho s_t - \gamma\Phi(n)\sigma(z)^2(1-\rho^2)\right] = 0.$$

At the interior optimum, the second-order condition implies $G_n(n^*, \rho; s_t) < 0$. Hence the implicit function theorem gives

$$\frac{\partial n^*}{\partial \rho} = -\frac{G_\rho}{G_n}.$$

Differentiating G with respect to ρ gives

$$G_\rho = \beta\Phi'(n^*)\sigma(z)\left[s_t + 2\gamma\Phi(n^*)\sigma(z)\rho\right].$$

Since $\beta > 0$, $\Phi'(n^*) > 0$, $\sigma(z) > 0$, and $G_n < 0$, the sign of $\partial n^*/\partial\rho$ is the sign of

$$s_t + 2\gamma\Phi(n^*)\sigma(z)\rho,$$

which proves (20).

For the case with endogenous z, let

$$W(n, z; \rho, s_t) = u(m_t - n(q_t + z)) + \beta K(n, z; \rho, s_t).$$

The first-order conditions are

$$W_n(n^*, z^*; \rho, s_t) = 0, \qquad W_z(n^*, z^*; \rho, s_t) = 0.$$

Totally differentiating this system with respect to ρ yields

$$H\begin{pmatrix}\partial n^*/\partial\rho \\ \partial z^*/\partial\rho\end{pmatrix} + \begin{pmatrix}W_{n\rho} \\ W_{z\rho}\end{pmatrix} = 0,$$

where H is the Hessian of W with respect to (n, z). Direct differentiation gives

$$W_{n\rho} = \beta\Phi'(n^*)\sigma(z^*)\left[s_t + 2\gamma\Phi(n^*)\sigma(z^*)\rho\right],$$

and

$$W_{z\rho} = \beta\Phi(n^*)\sigma_z(z^*)\left[s_t + 2\gamma\Phi(n^*)\sigma(z^*)\rho\right].$$

Solving the linear system gives (21). □

A.5 Proof of Proposition 4

Proof. Wherever the worst-case belief is interior, (22) gives

$$\rho^w(n, z, s_t) = -\frac{s_t}{\gamma\Phi(n)\sigma(z)}.$$

Substituting this value into (18) gives the minimized continuation value (26),

$$\min_{\rho} K(n, z; \rho, s_t) = \Phi(n)\mu(z) - \frac{\gamma}{2}\Phi(n)^2\sigma(z)^2 - \frac{s_t^2}{2\gamma}.$$

The last term depends on the signal, but it does not depend on either choice variable, n or z.

Therefore, on any region of the choice space where the worst-case belief is interior, the parent's problem is the signal-free problem (23) plus the constant $-\beta s_t^2/(2\gamma)$. Adding a constant changes the level of value but not the maximizer.

The signal restriction (24) is equivalent to

$$\underline{\rho} < -\frac{s_t}{\gamma\Phi(n^0)\sigma(z^0)} < \overline{\rho}.$$

Thus the worst-case belief is interior at (n^0, z^0). By continuity, it remains interior in a neighborhood of (n^0, z^0). On that neighborhood, the objective is exactly the signal-free objective plus a constant. Since (n^0, z^0) is the unique local maximizer of the signal-free problem, it is also the unique local maximizer of the ambiguous-signal problem. Hence

$$n^*(s_t) = n^0, \qquad z^*(s_t) = z^0,$$

and the budget condition gives

$$c^*(s_t) = m_t - n^0(q_t + z^0).$$

The constant-shift argument extends globally only across plans whose worst-case belief is interior. For a plan whose worst case sits at a boundary $\hat{\rho} \in \{\underline{\rho}, \overline{\rho}\}$, completing the square in (18) gives

$$\min_{\rho \in [\underline{\rho}, \overline{\rho}]} K = \Phi(n)\mu(z) - \frac{\gamma}{2}\Phi(n)^2\sigma(z)^2 - \frac{s_t^2}{2\gamma} + \frac{\gamma}{2}\left[\Phi(n)\sigma(z)\hat{\rho} + \frac{s_t}{\gamma}\right]^2,$$

whose last term is plan-dependent and nonnegative, so such plans move with the signal relative to (n^0, z^0). If the worst-case belief is interior at every candidate optimizer, the signal-free objective is globally strictly concave on that set, and (n^0, z^0) weakly dominates all boundary-belief plans and $n = 0$, then (25) is the unique global optimum.

The final step checks the endpoints of (24). At an endpoint the worst-case belief at (n^0, z^0) no longer lies strictly inside the ambiguity interval. It sits exactly at one boundary $b \in \{\underline{\rho}, \overline{\rho}\}$, so (n^0, z^0) is the knife-edge case between the interior regime and the boundary regime. Write $A(n, z) = \Phi(n)\sigma(z)$ for the exposure of a plan and $A^0 = A(n^0, z^0)$. At the endpoint the signal satisfies $s_t = -\gamma b A^0$. Substituting this into the boundary formula above shows that the minimized continuation value on the boundary side equals its interior-side counterpart plus the extra term $\frac{\gamma b^2}{2}[A(n, z) - A^0]^2$. The extra term vanishes at (n^0, z^0) and is positive at any plan with a different exposure. It is a convex reward for moving away from (n^0, z^0), so it pushes against the concavity of the objective, and whether (n^0, z^0) survives as a local maximum is a second-order comparison. Let H_0 be the Hessian of the interior-regime objective at (n^0, z^0), which is negative definite by assumed strict concavity. After discounting by β, the extra term adds the Hessian $\beta\gamma b^2 \nabla A \nabla A'$, a positive semidefinite matrix of rank one. The plan (n^0, z^0) remains a strict local maximum at the endpoint exactly when concavity wins, that is, when $H_0 + \beta\gamma b^2 \nabla A \nabla A' \prec 0$. This is the strengthened second-order condition referenced in the proposition. In the one-dimensional closed form, with $\Phi(n) = n$ and z fixed, both matrices are scalars. There $H_0 = -\beta\gamma\sigma^2$ and the extra term adds $\beta\gamma b^2\sigma^2$, so the condition reads $-\beta\gamma\sigma^2(1 - b^2) < 0$. It holds automatically because $b^2 < 1$. □

A.6 Proof of Proposition 5

Proof. Write the inner objective of (32) as

$$W(n; \rho, \alpha, s_t) = y - (q - \alpha\tau)n + \beta\left[n(\mu + \sigma\rho s_t) - \frac{\gamma}{2}n^2\sigma^2(1 - \rho^2)\right].$$

Part (i). For any (n, ρ) with $n > 0$ and $\tau > 0$,

$$\frac{\partial W}{\partial \alpha} = \tau n \geq 0,$$

so W is nondecreasing in α and the inner minimum over $\alpha \in [\underline{\alpha}, \overline{\alpha}]$ is attained at $\alpha^* = \underline{\alpha}$ for every plan and every ρ. (If $n = 0$ or $\tau = 0$, the objective does not depend on α and the selection $\alpha^* = \underline{\alpha}$ is without loss.) Because α enters W only through the additive term $\alpha\tau n$, which does not depend on ρ, the joint minimization over (ρ, α) separates,

$$\min_{\rho,\alpha} W(n; \rho, \alpha, s_t) = \min_{\rho \in [\underline{\rho}, \overline{\rho}]} W\Big(n; \rho, \underline{\alpha}, s_t\Big).$$

The resulting problem is exactly the closed-form problem of Section 3.5 with the child cost $q_e(\tau) = q - \underline{\alpha}\tau$ in place of q. All results of Sections 3.1–3.5 therefore apply verbatim with this substitution. In particular, (28) becomes (33).

Part (ii). Differentiating (33) with respect to τ,

$$\frac{\partial n^0}{\partial \tau} = \frac{\underline{\alpha}}{\beta\gamma\sigma^2}.$$

Part (iii). By Part (i), the household's objective under the ambiguous program $(\tau, [\underline{\alpha}, \overline{\alpha}])$ coincides, for every plan n and every signal s_t, with the objective under a known pass-through $\alpha \equiv 1$ and statutory size $\underline{\alpha}\tau$, since both produce the same effective cost $q - \underline{\alpha}\tau$. Identical objectives imply identical argmax correspondences, hence identical policies. Matching a legible program of size τ^* requires $\underline{\alpha}\tau = \tau^*$, i.e. $\tau = \tau^*/\underline{\alpha}$.

Part (iv). By Part (i) and the definition of λ in Section 3.5,

$$\lambda(\tau) = \frac{\mu - (q - \underline{\alpha}\tau)/\beta}{\sigma}, \qquad \frac{\partial \lambda}{\partial \tau} = \frac{\underline{\alpha}}{\beta\sigma} > 0$$

whenever $\underline{\alpha} > 0$, and the width of the inertia region is $\lambda(\tau)(\overline{\rho} - \underline{\rho})$ by (29). The illegible-program limit discussed after Proposition 5 follows by taking $\underline{\alpha} \to 0$ in Part (ii) and noting that the realized outlay per household is $\alpha_0 \tau n^0 > 0$ whenever $\alpha_0 > 0$ and $n^0 > 0$. $\square$

A.7 Dynamic consistency of the composed criterion

This appendix expands on the discussion following Assumption 2. Dynamic consistency asks whether the parent would ever want to abandon a plan midway. A criterion is dynamically consistent when the plan that looks best at date t, viewed from the whole future, still looks best at date $t + 1$ once the recursion is rolled forward one step. The criterion in (6) stacks two layers on top of the expectation, a worst case over ρ and a CARA certainty equivalent, and each layer raises this question separately.

The maxmin layer is the familiar case. Worst-case criteria are vulnerable to a specific failure. If nature must commit to one ρ for all time, the parent at date 0 guards against a single worst ρ for the whole future, while the parent at date 1 re-solves against a possibly different worst ρ, and the two plans can disagree. Epstein and Schneider (2003) show that this failure disappears when the set of priors over paths is rectangular, meaning that nature can pick the worst one-step-ahead law afresh at every date and history. Assumption 2 imposes exactly this structure. The path-prior set is built from arbitrary history-by-history selections $\rho_t(h_t) \in [\underline{\rho}, \overline{\rho}]$, so whatever conditional law the date 1 parent fears was already available to nature in the date 0 parent's problem. The two problems therefore share the same worst case, and the maxmin layer is dynamically consistent by the Epstein–Schneider theorem.

The CARA layer requires a different argument. The certainty equivalent $-\frac{1}{\gamma}\log \mathbb{E}[e^{-\gamma Y}]$ is a nonlinear transformation of the expectation, and the Epstein–Schneider theorem covers only the maxmin layer, not a nonlinear layer composed with it. No general theorem delivers consistency of the composition. We therefore follow the standard route in recursive robust control (Hansen and Sargent, 2008) and define preferences directly by the recursion (6). On this reading the parent does not first form a criterion over entire future paths and then check that the recursion reproduces it. The recursion itself is the primitive. Each generation evaluates its choice by current utility plus the composed adjustment of the next generation's value, and the question of a conflicting path-based evaluation never arises because no path-based evaluation is defined. Dynamic consistency holds by construction.

The entropy dual of Appendix A.8 gives the same conclusion a game interpretation. Each period the parent moves, and nature responds twice, once by selecting ρ_t from the interval and once by distorting the conditional law subject to an entropy penalty scaled by $1/\gamma$. Both of nature's moves are one-step-ahead moves inside the period, so the game is played afresh each generation and no commitment across periods is ever assumed or needed.

A.8 Properties of the distorted measure

This appendix defines the risk-sensitive distorted measure Q_t^* of Section 2 and records two of its properties. The measure reweights the worst-case transition law $P_{\rho_t^*,z_t^*}$ with Radon–Nikodym derivative

$$\frac{\mathrm{d}Q_t^*}{\mathrm{d}P_{\rho_t^*,z_t^*}} = \frac{\exp\{-\gamma\Phi(n_t^*)V(x_{t+1},s_{t+1})\}}{\mathbb{E}_{\rho_t^*,z_t^*}\left[\exp\{-\gamma\Phi(n_t^*)V(x_{t+1},s_{t+1})\} \mid x_t, s_t\right]}.$$

First, the likelihood reweighting is monotone in continuation value. For two successor states ω_1 and ω_2,

$$\frac{(\mathrm{d}Q_t^*/\mathrm{d}P_{\rho_t^*,z_t^*})(\omega_1)}{(\mathrm{d}Q_t^*/\mathrm{d}P_{\rho_t^*,z_t^*})(\omega_2)} = \exp\left\{-\gamma\Phi(n_t^*)\Big[V(\omega_1) - V(\omega_2)\Big]\right\},$$

so if $V(\omega_1) < V(\omega_2)$, then ω_1 receives relatively more weight under Q_t^* than under $P_{\rho_t^*, z_t^*}$. The distorted measure shifts probability mass toward low-value future states, more aggressively the larger is γ, and collapses back to $P_{\rho_t^*, z_t^*}$ as $\gamma \to 0$.

Second, the CARA certainty equivalent admits the entropy representation familiar from risk-sensitive and robust control,

$$-\frac{1}{\gamma} \log \mathbb{E}_{P_{\rho_t^*, z_t^*}} \left[e^{-\gamma Y_{t+1}} \mid x_t, s_t \right] = \min_{Q \ll P_{\rho_t^*, z_t^*}} \left\{ \mathbb{E}_Q[Y_{t+1} \mid x_t, s_t] + \frac{1}{\gamma} D(Q \| P_{\rho_t^*, z_t^*}) \right\},$$

where $Y_{t+1} = \Phi(n_t^*) V(x_{t+1}, s_{t+1})$ and D denotes relative entropy. The minimizer of this entropy-penalized problem is exactly Q_t^*. The representation shows in what sense Q_t^* acts like a pessimistic belief. It puts extra weight on bad future states, but the parent is not simply maximizing expected utility under Q_t^*. The distorted measure is generated by the risk-sensitive problem itself, changes with the chosen policy, and is disciplined by the entropy cost that limits how far it can move from the baseline worst-case law.

A.9 Risk-neutral and ambiguity-neutral limits

The model differs from the standard Barro–Becker problem in two ways. First, the parent is ambiguity averse. She evaluates each plan under the worst admissible signal-informativeness parameter $\rho \in [\underline{\rho}, \overline{\rho}]$. Second, she is risk sensitive. Future continuation values enter through the CARA certainty equivalent rather than through a standard expectation. This appendix separates the two forces explicitly (see Remark 2 in the main text).

Start by shutting down risk sensitivity. By Lemma 1, $\mathcal{C}_{\gamma,\rho,z}(Y) = \mathbb{E}_{\rho,z}[Y] - \frac{\gamma}{2} \operatorname{Var}_{\rho,z}(Y) + o(\gamma)$, so as $\gamma \to 0$ the certainty equivalent converges to a conditional expectation. Because $\Phi(n_t)$ is chosen at time t and is nonnegative, it can then be taken outside both the expectation and the minimization, and the Bellman equation becomes

$$V(x_t, s_t) = \max_{c_t, n_t, z_t} \left\{ u(c_t) + \beta \Phi(n_t) \min_{\rho \in [\underline{\rho}, \overline{\rho}]} \mathbb{E}_{\rho, z_t} \left[V(x_{t+1}, s_{t+1}) \mid x_t, s_t \right] \right\}. \tag{A.1}$$

Equation (A.1) is the risk-neutral robust Barro–Becker problem. Fertility still scales the value of descendants through $\Phi(n_t)$, but the continuation value is evaluated under the worst admissible signal-informativeness parameter.

Now shut down ambiguity instead. If the ambiguity set is a singleton, $\underline{\rho} = \overline{\rho} = \rho$, the inner minimization disappears and the problem is ambiguity-neutral but still risk sensitive. The parent uses a single transition law, indexed by ρ, and evaluates future outcomes through the CARA certainty equivalent in (6). Finally, if both features are switched off, that is $\gamma \to 0$ and $\underline{\rho} = \overline{\rho} = \rho$, the problem collapses to the standard Barro–Becker recursive formulation

$$V(x_t, s_t) = \max_{c_t, n_t, z_t} \left\{ u(c_t) + \beta\Phi(n_t)\mathbb{E}\left[V(x_{t+1}, s_{t+1}) \mid x_t, s_t\right] \right\}. \tag{A.2}$$

Thus the standard model is recovered only when the parent uses a single probability law and evaluates continuation value by an ordinary expectation.

B Robustness of the Inertia Result

This appendix states and proves the robustness results summarized in Section 3.7.

B.1 Smooth-ambiguity sketch

Let $I = [\underline{\rho}, \overline{\rho}]$ and let π be a probability measure on I. Under exponential smooth ambiguity, the parent evaluates a plan $\theta = (n, z)$ by

$$V_\eta(\theta, s) = -\frac{1}{\eta} \log \mathbb{E}_\pi \left[e^{-\eta K(\theta, \rho, s)} \right], \qquad \eta > 0,$$

where K is the LQ certainty equivalent (18). The second-order belief π is held fixed as s varies.

We show that the maxmin inertia result is recovered as ambiguity aversion becomes large. Let S be a compact subset of the maxmin interior-belief region and let N be a compact neighborhood of the signal-free optimizer θ_0. Suppose that, on $N \times S$,

$$A(\theta) \equiv \Phi(n)\sigma(z)$$

is bounded away from zero and

$$\rho^w(\theta, s) = -\frac{s}{\gamma A(\theta)}$$

stays a positive distance from the endpoints of I. Assume that π has a positive, sufficiently smooth density p near these minimizers and that the signal-free outer Hessian at θ_0 is nonsingular. Then there is, for all sufficiently large η, a local optimizer branch $\theta^*_\eta(s)$ near θ_0 satisfying

$$\sup_{s \in S} \left\| \frac{\partial \theta^*_\eta(s)}{\partial s} \right\| = O(\eta^{-1}).$$

Thus the exactly flat maxmin policy is the large-η limit of an attenuated smooth-ambiguity response.

To prove the claim, first define the ambiguity-adjusted decision weight

$$\frac{d\pi_{\eta,\theta,s}}{d\pi}(\rho) = \frac{e^{-\eta K(\theta,\rho,s)}}{\mathbb{E}_\pi\left[e^{-\eta K(\theta,\rho,s)}\right]}.$$

Since $K_s = A(\theta)\rho$, differentiation under the integral gives

$$\frac{\partial}{\partial s} \mathbb{E}_{\pi_{\eta,\theta,s}}[\rho] = -\eta A(\theta) \operatorname{Var}_{\pi_{\eta,\theta,s}}(\rho) \leq 0.$$

Hence a lower signal shifts weight toward larger values of ρ, the same direction as the maxmin selector $\rho^w(\theta, s)$.

On the interior-belief region, completing the square yields

$$K(\theta, \rho, s) = M(\theta, s) + \frac{\gamma A(\theta)^2}{2}\Big(\rho - \rho^w(\theta, s)\Big)^2,$$

where

$$M(\theta, s) = \min_{\rho \in I} K(\theta, \rho, s) = \Phi(n)\mu(z) - \frac{\gamma}{2}A(\theta)^2 - \frac{s^2}{2\gamma}.$$

Laplace's method, uniformly on $N \times S$ and through the derivatives used below, gives

$$V_\eta(\theta, s) = M(\theta, s) + \frac{\log \eta}{2\eta} + \frac{1}{2\eta}\log\left(\frac{\gamma A(\theta)^2}{2\pi}\right) - \frac{1}{\eta}\log p\Big(\rho^w(\theta, s)\Big) + o(\eta^{-1}).$$

This expansion first shows that V_η converges uniformly to the maxmin value M. More importantly, $M_{\theta s} \equiv 0$ and the other s-dependent term is multiplied by η^{-1}, so

$$V_{\eta,\theta s} = -\frac{1}{\eta}\partial_{\theta s}\log p\Big(\rho^w(\theta, s)\Big) + o(\eta^{-1}) = O(\eta^{-1}).$$

For the outer objective

$$F_\eta(\theta, s) = u\Big(m_t - n(q_t + z)\Big) + \beta V_\eta(\theta, s),$$

the implicit function theorem supplies the nearby optimizer branch and a uniformly bounded inverse Hessian. Differentiating its first-order condition therefore gives

$$\frac{\partial \theta^*_\eta}{\partial s} = -F^{-1}_{\eta,\theta\theta}\beta V_{\eta,\theta s} = O(\eta^{-1}),$$

which proves the stated attenuation result. This is a local result on compact subsets of the interior-belief region; it does not impose a global shape on the finite-η policy.

B.2 Discrete fertility

Because fertility is an integer, a small change in the signal need not change the chosen family size. Ambiguity gives a stronger result. Compare two adjacent positive family sizes, n and $n+1$. When both plans have an interior worst-case belief, the signal changes their values by the same amount. Their value difference is therefore constant, so the parent cannot switch between them in that region. This argument does not apply to the childless plan: choosing $n = 0$ gives no continuation value and hence no common signal penalty.

Let $y \geq 0$. The effective cost per child is $q_e > 0$: it equals q without a subsidy and

$q - \underline{\alpha}\tau$ with the policy. The feasible family sizes are

$$\mathcal{N}_y = \left\{0, 1, \dots, \left\lfloor \frac{y}{q_e} \right\rfloor \right\}.$$

For $n \in \mathcal{N}_y$, write total robust value as

$$W_n(s) = y - q_e n + \beta \min_{\rho \in [\underline{\rho}, \overline{\rho}]} \left\{ n(\mu + \sigma \rho s) - \frac{\gamma}{2} n^2 \sigma^2 (1 - \rho^2) \right\},$$

with $W_0(s) = y$.

Proposition B.2.1 (Discrete fertility)**.** *Assume $0 \leq \underline{\rho} < \overline{\rho} < 1$ and $\beta, \gamma, \sigma, q_e > 0$. Then the following statements hold.*

(i) Fix a signal s and allow n to vary continuously over $[0, y/q_e]$. The resulting value function is strictly concave. Thus there is one optimal integer, except possibly when two adjacent integers are tied.

(ii) Let $n \geq 1$ and suppose $n, n+1 \in \mathcal{N}_y$. Define

$$C_n = \Big[-\gamma \sigma \overline{\rho} n, -\gamma \sigma \underline{\rho}(n+1) \Big].$$

The interval C_n has positive length if and only if

$$\frac{\underline{\rho}}{\overline{\rho}} < \frac{n}{n+1}.$$

On all of C_n,

$$W_{n+1}(s) - W_n(s) = c_n \equiv \beta \left[\mu - \frac{\gamma}{2} \sigma^2 (2n+1) \right] - q_e. \tag{B.1}$$

Inside C_n, both plans have an interior worst-case belief. At an endpoint, one belief reaches a bound, but the value formula remains valid by continuity. If $c_n \neq 0$, the same plan is preferred at every signal in C_n. If $c_n = 0$, the two plans are tied throughout C_n.

(iii) The difference $\Delta_n(s) = W_{n+1}(s) - W_n(s)$ never decreases as s rises. If $\underline{\rho} > 0$ and C_n has positive length, it rises strictly outside C_n. Under these conditions, the two plans have a unique switching point when $c_n \neq 0$. It lies to the left of C_n if $c_n > 0$ and to the right if $c_n < 0$. If $c_n = 0$, the plans are tied on all of C_n.

Proof. For $n \geq 1$, completing the square in ρ gives

$$
\begin{aligned}
&n(\mu + \sigma\rho s) - \frac{\gamma}{2}n^2\sigma^2(1-\rho^2) \\
&\qquad = n\mu - \frac{\gamma}{2}n^2\sigma^2 - \frac{s^2}{2\gamma} + \frac{\gamma}{2}n^2\sigma^2\left(\rho + \frac{s}{\gamma n\sigma}\right)^2.
\end{aligned}
\tag{B.2}
$$

Thus the unique worst-case belief is

$$
\rho_n^w(s) = \mathrm{Proj}_{[\underline{\rho},\overline{\rho}]}\left(-\frac{s}{\gamma n\sigma}\right),
$$

and the minimized continuation value is

$$
\underline{K}_n(s) = \begin{cases} n\mu + n\sigma\overline{\rho}s - \dfrac{\gamma}{2}n^2\sigma^2(1-\overline{\rho}^2), & s \leq -\gamma n\sigma\overline{\rho}, \\[2ex] n\mu - \dfrac{\gamma}{2}n^2\sigma^2 - \dfrac{s^2}{2\gamma}, & -\gamma n\sigma\overline{\rho} \leq s \leq -\gamma n\sigma\underline{\rho}, \\[2ex] n\mu + n\sigma\underline{\rho}s - \dfrac{\gamma}{2}n^2\sigma^2(1-\underline{\rho}^2), & s \geq -\gamma n\sigma\underline{\rho}. \end{cases}
\tag{B.3}
$$

The three pieces meet with the same first derivative. For fixed s, the second derivative of the continuous value function is one of

$$
-\beta\gamma\sigma^2(1-\overline{\rho}^2), \qquad -\beta\gamma\sigma^2, \qquad -\beta\gamma\sigma^2(1-\underline{\rho}^2).
$$

Each expression is negative, so the continuous value function is strictly concave. It can therefore have at most two optimal integers, and two can be optimal only if they are adjacent. This proves part (i).

For plan n, the unconstrained minimizer belongs to $[\underline{\rho},\overline{\rho}]$ exactly when

$$
-\gamma\sigma\overline{\rho}n \leq s \leq -\gamma\sigma\underline{\rho}n.
$$

The same condition for plan $n+1$ gives another interval. Their overlap is C_n, which has positive length exactly when $\overline{\rho}n > \underline{\rho}(n+1)$. On C_n, the middle line of (B.3) applies to both plans, with the endpoint values included by continuity. Both values contain the term $-\beta s^2/(2\gamma)$, so it cancels when the plans are compared:

$$
\begin{aligned}
W_{n+1}(s) - W_n(s) &= -q_e + \beta\left[\mu - \frac{\gamma}{2}\sigma^2\Big((n+1)^2 - n^2\Big)\right] \\
&= \beta\left[\mu - \frac{\gamma}{2}\sigma^2(2n+1)\right] - q_e.
\end{aligned}
$$

This proves part (ii).

For part (iii), the envelope theorem gives

$$\frac{\partial W_n(s)}{\partial s} = \beta n \sigma \rho_n^w(s), \qquad n \geq 1.$$

Let $x = -s/(\gamma\sigma)$. Rescaling the projection gives

$$n\rho_n^w(s) = \mathrm{Proj}_{[n\underline{\rho}, n\overline{\rho}]}(x).$$

It follows that

$$\begin{aligned}\Delta_n'(s) = \beta\sigma \Big[&\mathrm{Proj}_{[(n+1)\underline{\rho},(n+1)\overline{\rho}]}(x) \\ &- \mathrm{Proj}_{[n\underline{\rho}, n\overline{\rho}]}(x) \Big] \geq 0.\end{aligned} \tag{B.4}$$

The first projection interval has larger lower and upper endpoints than the second, so the difference is nonnegative. If $\underline{\rho} > 0$ and the intervals overlap, the difference is zero exactly when $x \in [(n+1)\underline{\rho}, n\overline{\rho}]$, or equivalently when $s \in C_n$. Also,

$$\lim_{s \to -\infty} \Delta_n(s) = -\infty, \qquad \lim_{s \to +\infty} \Delta_n(s) = +\infty$$

when $\underline{\rho} > 0$. Thus $\Delta_n(s)$ moves continuously from negative to positive values, rises strictly outside C_n, and equals c_n on C_n. This proves the threshold statements. □

The condition $n \geq 1$ matters. Because $W_0(s) = y$, the childless plan has no common signal penalty. When the one-child plan has an interior worst-case belief,

$$W_1(s) - W_0(s) = \beta \left[\mu - \frac{\gamma}{2}\sigma^2 \right] - q_e - \frac{\beta s^2}{2\gamma}. \tag{B.5}$$

This difference depends on s and may equal zero inside the one-child interior region. Thus Proposition B.2.1 rules out switching inside a common-interior region only when both family sizes are positive.

The result compares one pair of plans; it does not by itself identify the global optimum. Another family size may become relevant inside C_n. But when n and $n+1$ are the two relevant choices, (B.1) fixes their ranking on C_n, even though the value of each plan changes with the signal. Welfare can therefore move while the chosen family size does not.

For comparison, with known $\rho \in (0,1)$ the adjacent value difference is

$$-q_e + \beta \left[\mu + \sigma\rho s - \frac{\gamma}{2}\sigma^2(1-\rho^2)(2n+1) \right],$$

so its unique threshold is

$$s_n(\rho) = \frac{q_e/\beta - \mu + \frac{\gamma}{2}\sigma^2(1-\rho^2)(2n+1)}{\sigma\rho},$$

and the width of an unconstrained interior band is

$$s_n(\rho) - s_{n-1}(\rho) = \gamma\sigma\frac{1-\rho^2}{\rho}.$$

This width falls as ρ rises. In the ambiguity model, if both plans choose the upper bound $\overline{\rho}$, their switching point is the same as in the known-$\overline{\rho}$ model. The same is true at the lower bound $\underline{\rho}$. Ambiguity therefore does not make every integer-choice band wider than every known-ρ benchmark. In a transition region, one plan may choose a boundary belief while the other still chooses an interior belief, so no single known-ρ formula describes their comparison.

References

Al Jazeera. Fears for future as South Korea's fertility rate drops again. News report, 28 February, https://www.aljazeera.com/news/2024/2/28/fears-for-future-as-south-koreas-fertility-rate-drops-again, 2024.

Robert J. Barro and Gary S. Becker. Fertility choice in a model of economic growth. *Econometrica*, 57(2):481–501, 1989. doi: 10.2307/1912563.

Stefan Bauernschuster, Timo Hener, and Helmut Rainer. Children of a (policy) revolution: The introduction of universal child care and its effect on fertility. *Journal of the European Economic Association*, 14(4):975–1005, 2016. doi: 10.1111/jeea.12158.

Gary S. Becker and Robert J. Barro. A reformulation of the economic theory of fertility. *The Quarterly Journal of Economics*, 103(1):1–25, 1988. doi: 10.2307/1882640.

Gary S. Becker and H. Gregg Lewis. On the interaction between the quantity and quality of children. *Journal of Political Economy*, 81(2, Part 2):S279–S288, 1973. doi: 10.1086/260166.

Judson A. Caskey. Information in equity markets with ambiguity-averse investors. *Review of Financial Studies*, 22(9):3595–3627, 2009. doi: 10.1093/rfs/hhn062.

Alma Cohen, Rajeev Dehejia, and Dmitri Romanov. Financial incentives and fertility. *Review of Economics and Statistics*, 95(1):1–20, 2013. doi: 10.1162/REST_a_00342.

John M. Danskin. The theory of max-min, with applications. *SIAM Journal on Applied Mathematics*, 14(4):641–664, 1966. doi: 10.1137/0114053.

David de la Croix and Matthias Doepke. Inequality and growth: Why differential fertility matters. *American Economic Review*, 93(4):1091–1113, 2003. doi: 10.1257/000282803769206214.

Matthias Doepke and Fabrizio Zilibotti. The macroeconomics of child labor regulation. *American Economic Review*, 95(5):1492–1524, 2005. doi: 10.1257/000282805775014425.

Matthias Doepke, Anne Hannusch, Fabian Kindermann, and Michèle Tertilt. The economics of fertility: A new era. In Shelly Lundberg and Alessandra Voena, editors, *Handbook of the Economics of the Family*, volume 1, pages 151–254. North-Holland, Amsterdam, 2023. doi: 10.1016/bs.hefam.2023.01.003.

Larry G. Epstein and Martin Schneider. Recursive multiple-priors. *Journal of Economic Theory*, 113(1):1–31, 2003. doi: 10.1016/S0022-0531(03)00097-8.

Larry G. Epstein and Martin Schneider. Learning under ambiguity. *The Review of Economic Studies*, 74(4):1275–1303, 2007. doi: 10.1111/j.1467-937X.2007.00464.x.

Larry G. Epstein and Martin Schneider. Ambiguity, information quality, and asset pricing. *The Journal of Finance*, 63(1):197–228, 2008. doi: 10.1111/j.1540-6261.2008.01314.x.

Oded Galor and David N. Weil. Population, technology, and growth: From malthusian stagnation to the demographic transition and beyond. *American Economic Review*, 90 (4):806–828, 2000. doi: 10.1257/aer.90.4.806.

Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989. doi: 10.1016/0304-4068(89) 90018-9.

Libertad González. The effect of a universal child benefit on conceptions, abortions, and early maternal labor supply. *American Economic Journal: Economic Policy*, 5(3): 160–188, 2013. doi: 10.1257/pol.5.3.160.

Raffaele Guetto, Giacomo Bazzani, and Daniele Vignoli. Narratives of the future and fertility decision-making in uncertain times: An application to the COVID-19 pandemic. *Vienna Yearbook of Population Research*, 20:223–260, 2022. doi: 10.1553/populationyearbook2022.res1.6.

Lars Peter Hansen and Thomas J. Sargent. Discounted linear exponential quadratic Gaussian control. *IEEE Transactions on Automatic Control*, 40(5):968–971, 1995. doi: 10.1109/9.384242.

Lars Peter Hansen and Thomas J. Sargent. Robust control and model uncertainty. *American Economic Review*, 91(2):60–66, 2001. doi: 10.1257/aer.91.2.60.

Lars Peter Hansen and Thomas J. Sargent. *Robustness*. Princeton University Press, Princeton, NJ, 2008. doi: 10.1515/9781400829385.

Philipp K. Illeditsch, Jayant V. Ganguli, and Scott Condie. Information inertia. *Journal of Finance*, 76(1):443–479, 2021. doi: 10.1111/jofi.12979.

Cosmin L. Ilut and Martin Schneider. Ambiguous business cycles. *American Economic Review*, 104(8):2368–2399, 2014. doi: 10.1257/aer.104.8.2368.

Nengjiu Ju and Jianjun Miao. Ambiguity, learning, and asset returns. *Econometrica*, 80 (2):559–591, 2012. doi: 10.3982/ECTA7618.

Wookun Kim. Baby bonus, fertility, and missing women. *Journal of Human Resources*, 2024. doi: 10.3368/jhr.0322-12208R3. Forthcoming.

Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji. A smooth model of decision making under ambiguity. *Econometrica*, 73(6):1849–1892, 2005. doi: 10.1111/j.1468-0262.2005.00640.x.

Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji. Recursive smooth ambiguity preferences. *Journal of Economic Theory*, 144(3):930–976, 2009. doi: 10.1016/j.jet.2008.10.007.

Heejong Kwag, Songhee Baek, and Myounggu Kang. Impact of social trust on young adults' childbearing intention: Why does South Korea show ultra-low fertility? *International Journal of Urban Sciences*, 28(4):585–610, 2024. doi: 10.1080/12265934.2024.2320905.

Rodolfo E. Manuelli and Ananth Seshadri. Explaining international fertility differences. *Quarterly Journal of Economics*, 124(2):771–807, 2009. doi: 10.1162/qjec.2009.124.2.771.

Paul Milgrom and Ilya Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, 2002. doi: 10.1111/1468-0262.00296.

Kevin Milligan. Subsidizing the stork: New evidence on tax incentives and fertility. *The Review of Economics and Statistics*, 87(3):539–555, 2005. doi: 10.1162/0034653054638382.

Priya Ranjan. Fertility behaviour under income uncertainty. *European Journal of Population*, 15(1):25–43, 1999. doi: 10.1023/A:1006106527618.

Anna Raute. Can financial incentives reduce the baby gap? evidence from a reform in maternity leave benefits. *Journal of Public Economics*, 169:203–222, 2019. doi: 10.1016/j.jpubeco.2018.07.010.

Assaf Razin and Uri Ben-Zion. An intergenerational model of population growth. *American Economic Review*, 65(5):923–933, 1975.

Kamila Sommer. Fertility choice in a life cycle model with idiosyncratic uninsurable earnings risk. *Journal of Monetary Economics*, 83:27–38, 2016. doi: 10.1016/j.jmoneco.2016.08.002.

Statistics Korea. Birth statistics in 2023 (provisional results). Statistical Press Release, Daejeon, 2024.

Thomas D. Tallarini. Risk-sensitive real business cycles. *Journal of Monetary Economics*, 45(3):507–532, 2000. doi: 10.1016/S0304-3932(00)00012-X.

Daniele Vignoli, Raffaele Guetto, Giacomo Bazzani, Elena Pirani, and Alessandra Minello. A reflection on economic uncertainty and fertility in Europe: The narrative framework. *Genus*, 76:28, 2020. doi: 10.1186/s41118-020-00094-3.