Parental Guidance and Supervised Learning

Alessandro Lizzeri Marciano Siniscalchi*

May 11, 2007

Abstract

We propose a simple theoretical model of supervised learning that is potentially useful to interpret a number of empirical phenomena relevant to the *nature-nurture* debate. The model captures a basic trade-off between *sheltering* the child from the consequences of his mistakes, and allowing him to *learn from experience*. We characterize the optimal parenting policy and its comparative-statics properties. We then show that key features of the optimal policy can be useful to interpret provocative findings from *behavioral genetics*.

1 Introduction

The nature-nurture debate has been one of the most controversial in the social sciences.[1] Since the 1980s, the literature on *behavioral genetics* has presented data that is broadly interpreted as tilting the balance strongly toward nature.[2] Some go so far as to say that this literature provides evidence that parents have very little (or no) effect on a variety of measures of the personality of their children.[3] At the same time, in wealthy countries, the amount of resources invested by parents in their children has reached all–time highs. How can we reconcile the enormous resources invested in children with evidence that parenting explains an extremely small fraction of the variation in the characteristics of children? This paper provides a model of optimal parenting that is consistent with the data from behavioral genetics, but suggests a very different interpretation: in our model, parents can

*We thank Daron Acemoglu, Joe Altonji and Greg Duncan. Financial support from NSF Grants SES-0452317 (Lizzeri) and SES-0453088 (Siniscalchi) is gratefully acknowledged. Aaron Hauptmann provided excellent research assistance.

[1]For some history on this debate, see Pinker [53].

[2]For an overview of this literature, see Plomin et al. [55] or Reiss et al [56]. Some of the findings of this literature have been popularized by Harris [44], [45], Pinker [53] and Ridley [58].

[3]See for instance Harris [44], [45], Pinker [53] and Ridley [58].

have significant effects on the characteristics of their children. The model also allows a detailed structural interpretation for some of the data in this literature, and for measures of heritability that are of interest more broadly.

The literature on the consequences of parenting is virtually unanimous in recognizing that parental support is essential for functional development in extreme situations. Harlow and coauthors[4] separated infant monkeys from their mothers; these subjects developed severe emotional and cognitive problems. The discovery of children in Romanian orphanages, who were raised with very little human contact, provided a tragic counterpart to these studies, leading to similar conclusions. These children were in the third to tenth percentile for physical growth, and "grossly delayed" in motor and mental development.[5]

However, views on the impact of differences within the 'normal' range of variation in parenting are sharply divided. A sizable literature across the social sciences argues that these differences can have important effects. Specifically, socialization researchers provide numerous studies that attempt to relate variation in parenting styles with measures of adjustment.[6]

The behavioral genetics literature starts with a criticism of much of socialization research for failing to recognize that the correlations between parenting styles and children's outcomes could be due to the shared genes between parents and their biological children. Behavioral geneticists then attempt to isolate the effects of the genes through two complementary approaches. The more direct one is to compare twins raised together by their biological parents with twins raised apart by different adoptive families.[7] One of the problems with this approach is the limited number of twins raised apart. The other approach is to compare outcomes for children with varying degrees of genetic and environmental relatedness and employ statistical techniques to estimate the percentage of variation in personality traits that is explained by genetic factors. The findings of this literature paint a very different picture of the effects of parents. For instance, behavioral genetics studies consistently find that *twins reared together are just as similar as twins reared apart*. In fact, some studies even find that twins reared together are *less similar* than twins reared apart. These findings are interpreted according to a simple statistical model (known as the ACE model) as evidence that, once one controls for genetic factors, the impact of traditional measures of family environment on most personality traits is greatly diminished.[8]

[4]See for instance Griffin and Harlow [41] and Harlow and Zimmerman [42].

[5]For an interesting analysis of these children, see Chisholm [26].

[6]For a survey, see for instance Collins et al. [28] and Demo and Cox [32].

[7]There are several parallel projects that gather information on this front. The first, large-scale project of this kind was the Minnesota twin study: see Bouchard et al. [19].

[8]This broad conclusion is subject to two qualifications: first, there is evidence that improvements in the family environment has positive effects on children's cognitive ability, if one restricts attention to families of

We develop a model that permits a very different interpretation of the same findings. In our model, children's characteristics evolve through interactions with both the environment and their parents; we formalize this dynamic process by adapting a standard learning problem to allow for parental supervision.

Our model of *supervised learning* captures a basic trade-off between *sheltering* the child from the consequences of his mistakes, and allowing him to *learn from experience*. We characterize the optimal parenting policy and its comparative-statics properties; we then use the key features of the model to tackle the nature-nurture debate [19, 54, 58]. In this respect, the central implications of our model follow from the *genetic mediation* of the optimal parenting policy. Namely, parents tailor their behavior to the characteristics of their children, and to how they believe these characteristics are functional to the environment children face. Two important consequences are: (i) on average, adoptive parents shelter their children more than biological parents, implying that genes matter more for outcomes of adoptive children; and (ii) parents behave more similarly with monozygotic (identical) twins than with dizygotic (fraternal) twins; therefore, the greater similarity of monozygotic twins is partly due to the concordance of parental intervention.

1.1 Methodology and Summary of Results

Our model has four key features:

- There are two agents, the child and the parent. The parent is solely interested in the child's welfare, and is active for T periods. The child is active for $L > T$ periods.
- In each period, the child must perform some task. The parent has better information than the child about the correct way to perform it.
- The child *learns by doing*: at the end of each period, he receives a signal about the quality of his performance. However, *learning is costly*: the child's per-period payoff is lower the worse his performance.
- The parent's actions simultaneously modify (typically increase) the child's per-period payoff and distort (typically bias) the child's signal about his performance.

One of the key features of our model is the *tradeoff between sheltering the child and allowing him to learn*. We assume that, to some extent, allowing the child to learn means

low socioeconomic status (Turkheimer et al. [68]). Second, and of more immediate relevance to the present paper, recent studies suggest the intriguing possibility that parental intervention may actually respond to specific, genetically-determined traits of the child, and thus reinforce or attenuate them (Reiss et al [56]). Thus, family environment may have significant effects, even though its impact is "genetically mediated."

that he must be allowed to make his own mistakes. However, mistakes are costly and the parent wishes to protect the child from making mistakes that are too severe. We study how this tradeoff is optimally resolved. We show that the optimal policy of the parent partially shelters the child, implying that learning may be slowed down by the presence of the parent; on the other hand, the child learns at a much smaller cost than if he were on his own. We also show that the optimal policy for the parent is linear in the child's bias at any date. This allows us to focus on one key parameter of the model: the intensity of parental intervention.

We then investigate the dynamics of this central parameter. We show that the intensity of intervention displays subtle time patterns. Specifically, for relatively low discount factors, the intensity of intervention is eventually decreasing. For high discount factors, however, the intensity of intervention may be increasing over time. Also, we consider the effects of several parameters of interest that have a natural interpretation. In particular, we study how parental intervention responds to the *ease of learning* how to perform the task, which may reflect both the child's natural abilities and the complexity of the environment. Our analysis uncovers countervailing "short-run" and "long-run" effects.

To account for the evidence from behavioral genetics, we develop an extension of our basic model, characterized by a heterogeneous population of parents and children. In this variant of our model, it is natural to assume that children's prior information reflect their genetic endowment, and parents' prior information reflect both their genes and their past experiences. Furthermore, different assumptions about the correlation of prior information can be employed to model different patterns of genetic relatedness and rearing conditions. Once these parameters have been chosen, our model generates specific patterns of correlation in the performance measures of siblings; these can then be compared with empirical findings from behavioral genetics. As noted above, our model turns out to match these findings, for a broad and easily interpretable set of parameter choices.

1.2 Additional Background Literature

Learning has been exhaustively investigated in theoretical models by economists, statisticians, and psychologists.[9] However, these studies typically abstract from the fact that learning takes place under the supervision of parents, caregivers, teachers, advisors, and other experts for a considerable fraction of an individual's life. The economics literature has developed several models of investment in child quality, starting at least with Becker [7]); a particularly relevant contribution is the recent paper by Cunha and Heckman [31],

[9] In economics there is a vast literature that studies learning from various points of view: from Bayesian learning to adaptive learning to fictitious play. On theory, see e.g. Fudenberg and Levine [37]. Camerer [23], chapter 6 discusses both theory and experiments.

which analyzes the optimal life-cyle profile of investments in children in a model with dynamic complementarities. This approach enables the authors to capture some important facts about child development.

An enormous literature in developmental psychology addresses the effects of parental care on the development of children.[10] Yet, formal modeling of *supervised learning*, i.e. the relation between parental behavior and its effects on children's learning processes, is almost absent. This is the focus of the present paper.

The behavioral-genetics literature typically focuses on personality traits and measures of cognitive achievements, and does not typically explore outcomes such as educational attainment or earnings. In the economics literature, the degree to which a child's home environment supports learning (as measured e.g. by how often the mother reads to her child, or whether she helps him learn numbers) has been shown in some studies to have significant effects on cognitive achievement.[11] Recent contributions by Sacerdote [62, 63], Bjorklund, Lindahl and Plug [16] and others are methodologically closer to the behavioral-genetics literature. For instance, Sacerdote [63] analyzes a sample of Korean children randomly assigned to American adoptive families; he finds that maternal education has a significant positive effect on the educational attainment of adopted children, but a much larger effect on that of biological children. Bjorklund et al. [16] analyze Swedish adoption data and report significant effects for both adoptive and biological parents. Moreover, they find evidence for a positive interaction effect between postbirth environment (nurture) and prebirth factors (nature): as we discuss in Sec. 3, this is consistent with our approach.

Parental care is essential for the survival of infants and young children; indeed, this is the case for all mammals and many other animal species. However, there are vast differences in 'parenting strategies' across human societies. The structure of human families and the relation of parents to their children has significantly evolved through the course of human evolution.[12] Even in present-day societies, anthropologists have documented considerable heterogeneity in the way parents interact with their children.[13] Developmental

[10]It is impossible to be exhaustive in providing references. See Shonkoff and Deborah Phillips eds. [64] for a recent overview of the field with respect to early childhood development.

[11]See e.g. Carneiro, Heckman, and Masterov [24] and Todd and Wolpin [67]. Also, a sizable literature investigates the effects of maternal employment on children's cognitive achievement. Results are mixed: some find that employment is detrimental (Baydar and Brooks-Gunn, [6]; Desai et. al. [33]; Belsky and Eggebeen [8], Bernal [12]), others that it is beneficial (Vandell and Ramanan [69]). See also the debate on the effects of family size and birth order (e.g., Black, Devereux, and Salvanes) [17].

[12]See Lancaster and Lancaster [51] for a discussion of the very different parent/child relation that must have existed in hunter-gatherer societies relative to that in societies based on agriculture, and to that of modern parents.

[13]For example, Lamport-Commons and Miller [50], and Richman *et al.* [57] describe parental practices among the Gusii Kenyians and Highland Mayans, and highlight stark differences as compared to behavior

psychologists have also provided a classification of parenting styles for Western societies.[14] Finally, as regards mammals and other animals, there is an enormous amount of cross-species variation in the degree to which parents invest resources in their young.[15]

2 "Hand-Holding"

2.1 The basic model

Agents and Horizon. The model features two agents, the child (he) and the parent (she). The child lives for $L > 1$ periods, whereas the parent is active (i.e., able to supervise the child) for $T < L$ periods.

Actions and Payoffs The child must perform a task in every period. The real number M represents the correct way to perform the task *on average*; however, the correct way to perform the task at time $t = 1, \ldots, L$ is represented by i.i.d. normal random variables $X_1, \ldots, X_L$; every X_t has a normal distribution, with mean M and precision p_X.

The parent's and the child's actions at time t are also real numbers, respectively denoted by $\bar{a}_t$ and $\bar{b}_t$. If, at time t, the parent chooses action $\bar{a}_t$, the child chooses action $\bar{b}_t$, and the correct way to perform the task is X_t, then the child incurs a loss of

$$(X_t + \bar{a}_t - \bar{b}_t)^2.$$

As we discuss below in greater detail, the parent knows M, the correct way to perform the task on average, and can also anticipate the child's choice $\bar{b}_t$: thus, the parent's action $\bar{a}_t$ is effectively a *correction for the child's average mistake.*

Both the parent and the child wish to minimize discounted expected losses, given their respective information; thus, the parent is altruistic, The parent discounts per-period losses at a rate $\delta \in (0, 1)$. The child's discount factor may or may not coincide with that of the parent; however, because of subsequent assumptions, this plays no role in the analysis.

in the typical Western family. Blurton-Jones [18] describes studies of two populations of foragers of the sub-Saharan savanna: the !Kung and the Hadza. He documents that the former appear to adopt a much more protective and close relation with their infants as compared to the latter. Interestingly, this author suggests that this difference in parenting styles may be partly related to the harsher conditions of the land inhabited by the !Kung, and that this also leads to large differences in fertility between the !Kung and the Hadza.

[14] For instance, authoritative parenting, which is characterized by high levels of warmth, support and control is thought to be more effective than either authoritarian (which lacks warmth) or permissive (which lacks control) styles of parenting. See for instance the seminal study by Baumrind [5], and also Collins et al. [28] and Demo and Cox [32] for reviews of this literature.

[15] See for example Clutton-Brock [27].

Information and Policies. The child does not know M, but has prior beliefs about it. Specifically, we assume that, from the child's point of view, M is normally distributed, with mean M_0 and precision p_0. Similarly, as far as the child is concerned, the correct way to perform the task at time t, namely X_t, has a normal distribution *conditional upon* M, with mean M and precision p_X. The child also assumes that $X_1, \ldots, X_L$ are conditionally independent.

The child chooses his action a_t at the beginning of each period t. Upon completing the task, he receives feedback about his performance; however, he cannot distinguish between the consequences of his own choice and those of his parent's intervention. We model this by assuming that, at the end of each period t, the child observes the sum $X_t + a_t$, but not its separate components X_t and a_t.

The parent knows M, and also observes the realization of X_t at the end of period t. Furthermore, the parent knows the child's prior.[16]

Now temporarily suppose that the child was facing a standard learning model, without a parent. Under the usual assumption that the child minimizes his discounted expected losses, his optimal choice at time t would be the conditional expectation of M (equivalently, X_t) given the prior history. In particular, under the above normality assumptions, this takes a particularly convenient form:

$$E[X_t | X_1 = x_1, \ldots, X_{t-1} = x_{t-1}] = \frac{p_0 M_0 + p_X \sum_{s=1}^{t-1} x_s}{p_0 + (t-1) p_X}. \tag{1}$$

If instead the child faces a supervised-learning model, his optimal choices depend in part upon his understanding of the parent's own intervention policy. Our baseline model assumes that *the child disregards the parent's influence on his learning environment. Formally, he acts as if* $\bar{a}_t = 0$. We think of this as an interesting polar case that is helpful as an initial step. We also pursue a "textbook equilibrium" approach in the Web Appendix; our main findings remain true in this alternative model.

As a consequence of this assumption, at any time $t \leq T+1$, after observing $x_1 + \bar{a}_1, \ldots, x_{t-1} + \bar{a}_{t-1}$, the child's optimal action is his conditional expectation of X_t:

$$M_t^a \equiv E[X_t | X_1 = x_1 + \bar{a}_1, \ldots, X_{t-1} = x_{t-1} + \bar{a}_{t-1}] = \frac{p_0 M_0 + p_X \sum_{s=1}^{t-1} (x_s + \bar{a}_s)}{p_0 + (t-1) p_X}. \tag{2}$$

A similar expression for the conditional expectation M_t^a applies to periods $t > T+1$, when the prior history includes dates at which the parent is not active: see the Appendix for details. Also notice for future reference that $M_0^a = M_0$.

[16]This is not particularly restrictive in this version of the model. If the parent does not know the value of M_0 but anticipates the form of the child's policy, he can learn M_0 in one period, provided he knows p_0. Otherwise, he can learn both M_0 and p_0 in two periods.

Comparing Eqs. (1) and (2) immediately shows that parental intervention distorts the child's learning process. On the other hand, parental intervention directly affects the child's per-period payoff; in particular, the expected time-t penalty conditional upon the parent's information has a simple "variance plus bias" representation:

$$E[(X_t + a_t - M^a_{t-1})^2 | X_1, \ldots, X_{t-1}, M] = p_X^{-1} + (M + a_t - M^a_{t-1})^2. \tag{3}$$

Thus, Eqs. (2) and (3) reflect the basic tradeoff in this model.

As in the standard learning models we build upon, our agents are Bayesian rational (i.e. they maximize expected utility). However, our model is set up so that the *child*'s learning problem is elementary, and its solution involves a simple adaptive rule. We are confident that our main findings can be generalized to a suitable class of non-Bayesian-rational adaptive learning rules. On the other hand, we realize that certain aspects of child development, and hence parenting, are probably best understood in a setting where computational, memory, or motor constraints (and the evolution thereof) play a central role. A model that reflects these considerations likely requires substantial departures from Bayesian learning. In any event, adopting a relatively "conservative" modeling approach has two advantages: first, we can rest assured that the decision-theoretic foundations of our basic models are sound and well-understood; second, and more importantly, insofar as the conclusions of our analysis deviate from the predictions of standard models, we can be sure that this can be attributed entirely to expert supervision.

Another limitation is that we do not allow the parent to "describe" or "demonstrate" how to perform the task at hand. We only model one communication channel between the parent and the child, namely the former's intervention in the latter's learning process. We certainly do **not** wish to suggest that, in actual parent-child interactions, this really is the only open communication channel, and in fact we briefly discuss the empirical consequences of one highly stylized model of communication in Section 3.[17] We do not even wish to suggest that we believe it to be the most "interesting." Rather, the present paper focuses on one specific channel that, as the preceding discussion suggests, is important in a variety of settings. Developing a model of supervised learning that incorporates more "direct" forms of teaching as well as imitation[18] is an intriguing direction for further research.

[17] Indeed, it is well known at least since the work of Bandura and coauthors [4, 3] that children can learn by imitating the behavior of others. Furthermore, a stream of literature in developmental psychology, starting in part with the work of Vygotsky [70], emphasizes the "social" aspects of child development. Similar considerations are true for animals (see e.g. Zentall and Galef [77].

[18] See for instance Schlag [66] who provides a bounded rationality model with imitation; see also Apesteguia et al. [1] for a treatment with theory and experiments.

2.2 Three benchmark parenting policies

Before we analyze the solution to the parent's problem, it is useful to consider three reference, or benchmark parenting policies.

Letting Go: $a_t = 0$. This is the simplest policy. Clearly, it does not induce any bias in the child's learning process.

Full Sheltering: $a_t = M^a_{t-1} - M$ (a.k.a. "The Italian Mom"). Since the parent knows the value of M, this policy minimizes the per-period loss at times $t = 1, \ldots, T$: this can be seen from Eq. (3). Intuitively, recall that the child's choice at time t is $b_t = M^a_{t-1}$, and the loss is $(X_t + a_t - b_t)$: thus, by chooosing $a_t = M^a_{t-1} - M$, the parent "shifts" the mean of X_t so that it coincides with the child's choice. In other words, the parent makes sure that the child "gets it right" on average. Of course, this has negative consequences in terms of learning: the child's belief that M^a_{t-1} is the mean of X_t is reinforced, no matter how close or distant from the true mean M it may be.

The Boot Camp: $a_t = \frac{p_0+(t-1)p_X}{p_X}(M - M^a_{t-1})$. This policy ensures that, at the end of time t (i.e. after observing X_t), the child's posterior M^a_t will be equal to M on average. Intuitively, we can think of this policy as exacerbating the loss to the child for an incorrect choice, thereby accelerating learning.

Thus, the present framework allows for a range of qualitatively very different parenting strategies. Moreover, the Full Sheltering and Boot Camp policies will turn out to be useful reference points to understand the main features of the optimal solution: Full Sheltering maximizes myopic payoffs, whereas the Boot Camp policy maximizes learning.

2.3 Characterization and key features of the Optimal Policy

We can now state our main characterization result:

Theorem 2.1 *The optimal action of the parent at time t is a linear function of the child's bias: $a_t = \gamma_t(M^a_{t-1} - M)$. The intensity of intervention γ_t is time-varying but deterministic, and lies between zero and one. Also, γ_t is decreasing in δ and L. Finally, γ_t is a weighted average of the intensities of intervention for the "Full Sheltering" and "Boot Camp" policies.*

A formal statement and proof of this and all other results are in the Appendix.

The key qualitative conclusion of Theorem 2.1 is the finding that optimal parenting entails *partial sheltering*: the intensity of intervention γ_t lies in $(0, 1)$. This finding plays a central role in our analysis of the evidence from behavioral genetics in Sec. 3.

As the discount factor δ increases, and/or the number of unsupervised periods $L - T$ increases (specifically, if L increases and T is held fixed), learning the correct value of

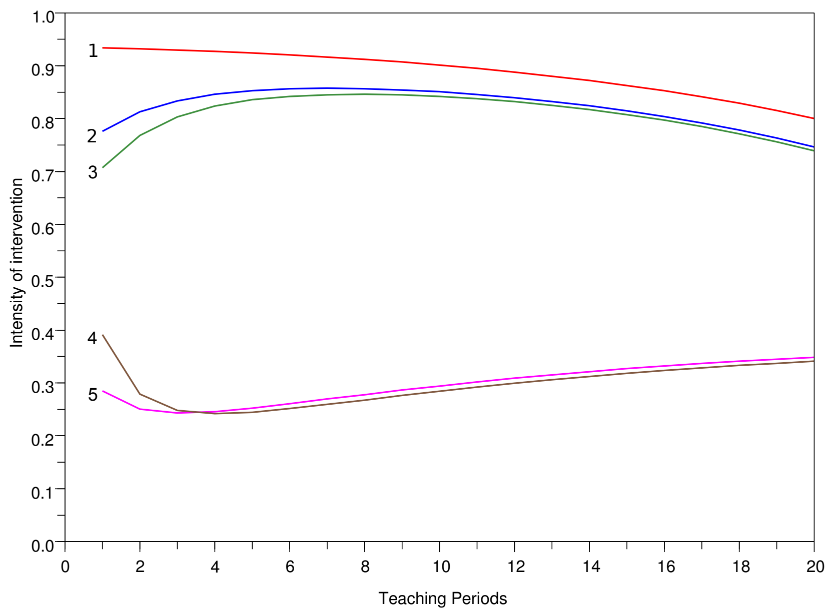

Figure 1: Intensity of intervention for different values of p_0 and p_X.

M becomes more important for the child. Theorem 2.1 confirms that, in this case, the intensity of intervention γ_t decreases.

Finally, the optimal policy is a combination of a parenting strategy that maximizes learning (the "Boot Camp") and one that maximizes myopic payoffs ("Full Sheltering").[19] In particular, $\gamma_t = \mu_t \gamma_t^{\text{FS}} + (1-\mu_t)\gamma_t^{\text{BC}}$, where $\gamma_t^{\text{FS}} = 1$, $\gamma_t^{\text{BC}} = -\frac{p_0+(t-1)p_X}{p_X}$, and $\mu_t \in (0,1)$. The weight μ_t placed on the "Boot Camp" intensity γ_t^{BC} can be shown to reflect the relative cost of the child's biases in the current and the following periods: μ_t (and hence γ_t) will be higher in periods when the cost of mistakes is high relative to the following period.

Illustration and Interpretation. Figure 1 depicts the intensity of intervention for a range of parameters, summarized in Table 1. The resulting patterns of parenting behavior are representative of what our model can generate.

In addition to the properties highlighted in Theorem 2.1, we draw attention to two key features related to, respectively, the time evolution of the intensity of intervention and its dependence on parameters related to ease of learning.

The **dynamics of the intensity of intervention** can be understood in terms of the decomposition of γ_t into a weighted average of the "Boot Camp" and "Full Sheltering"

[19] Thus, every parent is part drill sargeant and part Italian mom.

Label	Color	δ	p_0	p_X
1	Red	0.9	1	0.1
2	Blue	0.9	1	1
3	Green	0.9	0.1	1

Label	Color	δ	p_0	p_X
4	Brown	0.99	0.1	1
5	Magenta	0.99	1	1

Table 1: Parameters for the plots in Fig. 1. $L = 100$, $T = 20$.

intensities. It can be shown that, except possibly for the first few time periods, the weight μ_t placed on the "Full Sheltering" intensity γ_t^{FS} will be *decreasing* in t for δ relatively low and *increasing* in t for δ relatively high.

Consider first the case of relatively low discount factor; refer to the curves labeled 1,2 and 3 in Fig. 1. Recall first that, by Theorem 2.1, the intensity of intervention will be relatively high at any point in time. Now notice that γ_t^{FS} is constant; on the other hand, since the precision $p_0 + (t-1)p_X$ of the child's posterior at time t increases linearly with t, the coefficient γ_t^{BC} becomes more and more negative.[20] Since μ_t is eventually decreasing in t, the same will be true for γ_t. To interpret, note that the impact of observations on the child's posterior is greater early on (see Eq. (2)), so sheltering in later periods induces a smaller bias. We conclude that, if δ is low, the parent places less weight on reducing the child's bias than on minimizing current losses, and a *high, but eventually decreasing* level of sheltering is optimal.

A symmetric argument applies to the case of relatively high discount factor (curves 4 and 5 in Fig. 1), leading to *low, but eventually increasing* intensities of intervention. There is, however, an additional complication: even if the weight placed upon Full Sheltering increases, the intensity of intervention for the Boot Camp policy decreases linearly. Still, simulations suggest that the pattern displayed in Fig. 1 is prevalent.

Ease of learning is determined both by the child's ability to learn and the complexity of the environment. In our model, these are captured by the relative magnitude of the precisions p_0 and p_X. Refer to Eq. (1), which characterizes Bayesian updating of the estimated mean of X_t: if p_0 is high, or if p_X is low, the child places more weight on her prior M_0 than on observations, and hence learning occurs more slowly.

Our analysis identifies two effects of the ease of learning on the intensity of intervention. The first is straightforward: if learning occurs more slowly, the child benefits less from a reduced bias; thus, there is an incentive to provide more sheltering, i.e. *increase* γ_t, when learning is harder. We call this the "inertia" effect.

[20] Intuitively, the child's posterior is less affected by experience in later periods, and therefore a more substantial intervention is required in order to correct a given expected bias.

There is, however, a more subtle intertemporal effect, pushing in the opposite direction. If learning is harder, this will be the case not just today, but also in the future; in other words, the "cost" (continuation value) of the residual bias at the end of the current period is higher when learning is harder. Thus, there is an incentive to provide less sheltering, i.e. *decrease* γ_t, when learning is harder. We call this the "continuation value" effect.

We have verified (via numerical analysis) that the inertia effect dominates when the discount factor is low, whereas the continuation-value effect can prevail when δ is high.

3 Nature and Nurture: Interpreting evidence from behavioral genetics

3.1 A Population Model

In order to analyze the interaction between genetic effects and parenting, we embed our simple, two-agent supervised-learning model within a population framework where parents and children are heterogeneous.

We continue to assume that the correct way to perform a task is represented by the real number M. We also continue to assume that every child has a normal prior over M. The dimension of heterogeneity we explore is the simplest one to analyze in our model: we assume that a distribution of prior means M_0 in the population of children is given. Formally, we now treat M_0 as a random variable.

Symmetrically, we assume that parents do not observe M, and have a normal prior on it, with mean Z_0 and precision p_{Z0}. Furthermore, a distribution of the prior mean Z_0 in the population of parents is given: thus, Z_0 is also treated as a random variable.

Because M is the same for everyone in the population, the heterogeneity in priors, and in the corresponding posteriors, has payoff consequences, and can be interpreted as a a fitness measure: agents who are closer to the true mean, namely, those who are more correct on average, are better off because they make better decisions.

The exact distribution of M_0 and Z_0 is not important. We do, however, make a few assumptions relating the key uncertain quantities in the model.

First, to aid in the interpretation of this model, it is useful to think of the parent's prior beliefs $\mathrm{N}(Z_0, p_{Z0})$ as coming from (1) some prior belief that the parent held when she was born as a child, and (2) subsequent experience acquired as the parent was growing up. We also imagine that the distribution of the beliefs held by the parent as she was born as a child is the same as the distribution of the current child's beliefs. We shall elaborate on this point in Sec. 3.3.

We also assume that children's and parents' prior means are uncorrelated with the observations.

Whenever we consider more than one parent and/or more than one child, these assumptions will apply to each child-parent pair and to the observations made by the child in that pair. Observations made by different children will be assumed to be conditionally independent.

The correlations between two individuals' prior means play a crucial role in our analysis. Specifically, *these correlations reflect the genetic relatedness of the individuals under consideration*. Thus, for instance, the prior means of unrelated parents, or of parents and adoptive children, are uncorrelated; the prior means of dizygotic twins are perfectly correlated; and so on.[21]

An important feature of the linear-quadratic-Gaussian framework adopted here is that, even if parents do not know M, *the optimal policy has the same structure as in Theorem 2.1*; in particular, the intensity of intervention γ_t is the same as in the model of the previous section. The only difference is that, at each time t, M is replaced with its conditional expected value at that time, given the parent's prior mean Z_0 and the realizations of the signals $X_1, \ldots, X_t$.

3.2 Behavioral Genetics: the ACE model

As was noted in the Introduction, the literature on behavioral genetics (BG henceforth) empasizes the central role of an individual's genes in determining a variety of traits such as IQ or, more broadly, cognitive achievement. To interpret and provide a rationale for their conclusions, behavioral geneticists typically adopt the so-called "ACE model"; we shall now discuss a widely-used, if simple variant (cf. Plomin, DeFries, McClearn and McGuffin [55], p. 345 ff.).

The first step is to decompose the observable characteristic of interest, or *phenotype*, into a sum of three factors: the individual's *genotype*, or genetic endowment; the *shared environment*, corresponding to factors that affect siblings reared in the same family; and the *non-shared environment*, which captures idiosyncratic elements of the phenotype. The random variables corresponding to the phenotype and the three factors just described are commonly denoted by P, A, C and E respectively; the reference equation of the ACE

[21]We note that there may be some ambiguity as to the interpretation of the term 'parent.' We typically mean some aggregate of the two parents, which reflects the extent to which child-rearing responsibilities are shared within the family. This only matters when quantifying the genetic relatedness of a child and his biological 'parent.' However, we do not need to take a stand on the precise values here; all we need is that the child have substantial (but not perfect) genetic relatedness with his biological 'parent.'

model can then be written as

$$P = A + C + E. \tag{4}$$

The typical assumption in the ACE model is that the factors affecting a given individual's phenotype are *mutually independent*.[22] Consequently, the variance of P equals the sum of the variances of the three factors A, C and E.

It should be emphasized that the additive formulation in Eq. (4) does *not* follow from, or even suggest, an explicit biological–developmental "production function" whereby genetic and environmental inputs are transformed into behavioral outputs. Indeed, in the words of Goldberger [39], the factors A, C and E are themselves best viewed as "hypothetical constructs."

By way of contrast, our model suggests a linear relationship whose terms have a direct interpretation in our framework. Specifically, consider for simplicity a two-period version of the model in Sec. 2; the parent is active only in the first. We can interpret the child's posterior at time 1, i.e. M_1, as her phenotype. The form of the optimal parenting policy provided by Theorem 2.1 enable us to express the phenotype M_1 as a function of the primitive parameters of our model, as well as the intensity of intervention γ_1. Specifically, we write

$$\underbrace{M_1}_{P} = \underbrace{\frac{p_0}{p_0 + p_X} M_0}_{A} + \underbrace{\gamma_1 \frac{p_X}{p_0 + p_X}(M_0 - Z_0)}_{C} + \underbrace{\frac{p_X}{p_0 + p_X} X_1}_{E}, \tag{5}$$

where we have emphasized a possible mapping between the key quantities in our framework and the additive factors in the ACE model. The terms corresponding to A and E have been chosen to be purely genetic and, respectively, purely environmental and non-shared respectively; the remaining term captures the effects of parenting, and corresponds to the factor C in the ACE model.[23] Eq. (5) will provide the basis for the analysis in the following two subsections.

3.3 Evidence from Twin Studies

One simple approach to evaluate the relative importance of genes and common rearing in determining cognitive and behavioral traits entails computing the correlation between measured characteristics of the two members of a twin pair ("phenotypic correlation"), and comparing these correlations across distinct categories of twins that differ by genetic similarity and/or rearing.

[22] There are a few notable exceptions: e.g. Flynn...

[23] An alternative to the proposed mapping is to let $A = \frac{p_0 + \gamma_1 p_X}{p_0 + p_X} M_0$, leave E unchanged, and adjust C accordingly. This would only require minor alterations in the following discussion; the main conclusions of our analysis would be, of course, unchanged.

For instance, for most traits of interest, the phenotypic correlation for monozygotic (identical; MZ henceforth) twins reared *apart*, by different adoptive parents, is not significantly smaller than for MZ twins reared *together*, by their biological parents. Specifically, in most studies, the phenotypic correlation for MZ twins reared apart, denoted r^{MZA}, is at least 90% of the phenotypic correlation for MZ twins reared together, denoted r^{MZT}. In fact, for certain traits, phenotypic correlation is actually higher for twins raised apart than for twins raised together (Bouchard [19], Tab. 4).

Furthermore, the difference between r^{MZT} and r^{MZA} is considerably smaller than that between the phenotypic correlations for MZ and dizygotic (fraternal; DZ henceforth) twins reared together.

In the BG literature, these findings are interpreted as indicating that common rearing plays a minor role in determining the traits of interest. Intuitively, *if genotype and environmental factors are independent*, the phenotypic correlation between twins reared together is determined both by their common genetic endowment, and by their common rearing (the non-shared environment is, by definition, unique to each child); on the other hand, for twins reared apart, phenotypic correlation can only be driven by commonality in their genetic endowment. Thus, if r^{MZA} is not much smaller than r^{MZT}, the contribution of common rearing to phenotypic correlation must be small. Similarly, MZ and DZ twins reared together share the same rearing environment; taking the difference between their respective phenotypic correlations "cancels out" the common effects of parenting, so any remaining difference must be due to the fact that MZ twins have the same genetic endowment, whereas DZ twins only share 50% of the genes. Therefore, if $r^{MZT} - r^{MZA} < r^{MZT} - r^{DZT}$, reducing genetic commonality between twins has a greater effect on phenotypic correlation than rearing them in separate families. As shown in Appendix A.3, this intuition can be formalized in the ACE model—although it should be noted that the latter cannot explain the finding that $r^{MZA} > r^{MZT}$ for certain traits.

We shall now show that *these same patterns of phenotypic correlations arise in our model of supervised learning, precisely* <u>*because*</u> *of key features of the optimal parenting policy.* Therefore, the conclusion that family influence is limited is not necessarily warranted—in our model, the above correlational patterns are consistent with significant parental input.

3.4 Phenotypic Correlations under Supervised Learning

Note first that Eq. (5) may be used to compute phenotypic correlations for all twin categories mentioned in the preceding discussion. It is convenient to define $p = \frac{p_X}{p_0 + p_X}$, which corresponds to the precision of the observation X_t as a fraction of the precision of the child's posterior M_1. We also denote by v_0 and v_{Z0} the population variances of M_0 and

Z_0 respectively; by v_1 and v_{1a} the variances of M_1 for children reared by their biological parents and by adopted parents respectively (as we elaborate in §3.5 below, these will be different in our model); and by $v_X = \frac{1}{p_X}$ the variance of X_t. Finally, let r_0 denote the correlation between M_0 and Z_0 (for a child reared by her biological parents).

We show in Appendix A.4 that

$$r^{MZT} - r^{MZA} = [p\gamma_1]^2 \frac{v_{Z0}}{v_1} - r_0 \left\{ 2\frac{v_X + \gamma_1^2 v_{Z0}}{v_{1a}} [(1-p) + p\gamma_1] p^3 \gamma_1 \frac{\sqrt{v_0 v_{Z0}}}{v_1} \right\}, \quad (6)$$

$$r^{MZT} - r^{DZT} = \frac{1}{2}[(1-p) + p\gamma_1]^2 \frac{v_0}{v_1}. \quad (7)$$

It is useful to focus on Eq. (6). We emphasize a key qualitative feature: the correlation r_0 between the child's prior M_0 and her biological parent's initial mean Z_0 enters with a negative sign in Eq. (6), and hence *reduces the difference in phenotypic correlation between MZ twins reared together and reared apart*. As a result, the difference $r^{MZT} - r^{MZA}$ may well be very small, and even negative.

This effect is the main force that enables our model of supervised learning to generate the correlational patterns discussed above. Indeed, as we shall demonstrate momentarily, the above correlational patterns emerge for a broad range of sensible parameterizations. However, we first discuss the intuition behind the negative effect of r_0 on the difference in phenotypic correlation.

Each twin raised by adoptive parents is likely to be less similar to her parent than if she was raised by her biological parent. On average, this will lead adoptive parents to provide *more sheltering*. This is consistent with evidence from a variety of sources. For instance, Hoopes [46] finds that "adoptive mothers are more protective and careful with the children... adoptive mothers and fathers reported that they fostered more dependency than the biological fathers and mothers. The latter group admitted to greater feelings of irritability regarding their children, and the fathers tended to force independence, suppress affection, and accelerate development (p.23)." Furthermore, these more protective attitudes of adoptive parents "may have their effect on the children, who, at 5 years of age, were rated as a little less confident and less willing and attentive in task completion (p.27)." Warren [73] shows that "adoption significantly increases the likelihood of referral for psychiatric treatment, even after controlling for the fact that adoptees are significantly more likely to be referred when they display few problems."

But this means that non-shared environmental influences will have fewer opportunities to affect the twins' posterior, which, as a result, will be *more similar to their prior, and hence more similar to one another* on average. Conversely, twins raised by their biological parents will be subject to less sheltering, because their priors are positively correlated with their parents'. Hence, their non-shared experiences will have a greater role, and they will end up being less similar to one another.

Thus, differential sheltering by biological and adoptive parents provides a countervailing force to common rearing;[24] this is captured in Eq. (6) by the negative coefficient of r_0.

It should be clear that sheltering plays a crucial role in the mechanism just described. For instance, consider a variant of the model developed in this paper in which parents simply communicate their prior mean Z_0 to children, but do not intervene in their learning process. Such a model would differ from the reference ACE framework because shared environmental factors are correlated with the child's genes. Yet, such a model would be unable to match the empirical finding that, for certain traits, $r^{MZA} > r^{MZT}$. Furthermore, to account for the other correlational patterns described above, one would need to assume that the communication technology is considerably noisy (or, equivalently, that children do not place much weight on what their parents tell them). But, in a model of this kind, this would be tantamount to assuming directly that parents have limited influence on their children.

On the other hand, we can think of other models that generate analogous countervailing forces through different means. For instance, we could obtain similar effects if we assumed that intervention by parents is solely driven by a desire to have their children be similar to them.[25] Under this assumption, because the biological parents are more similar to their children than the adoptive parents, the latter intervene more intensely, thus generating similar effects to those discussed above.

We finally demonstrate in Figure 2 below that the correlational patterns described in the preceding subsection emerge for a broad range of parameters. In particular, to rule out implausible parameterizations,[26] we consider a "steady-state" version of our model. In every period, parents are explicitly modeled as individuals who were born at the beginning of the previous period with some prior beliefs, which they subsequently revised in light of observations made *under the supervision of their own parents*. This enables us to write the parent's initial mean Z_0 in a form analogous to Eq. (5), and compute its variance v_{Z0} and correlation r_0 with M_0 (for twins reared by their biological parents) as functions of the remaining parameters. Appendix A.4 provides the details.

[24] The fact that common rearing increases phenotypic correlation for twins reared together, and hence the difference between r^{MZT} and r^{MZA}, is captured by the positive coefficient on v_{Z0}, the variance of the parent's prior; see also Eq. (23) in Appendix A.4.

[25] Bisin and Verdier [13, 14] make such an assumption in a model of cultural transmission.

[26] For instance, we wish to rule out the possibility that Z_0 may be non-random, because this would be inconsistent with the assumption that this quantity is itself determined by the parent's prior at "time -1" and subsequent (supervised) learning.

We set $\delta = (0.95)^{20}$, intuitively suggesting a 95% yearly discount factor and a teaching period lasting for 20 years (higher values of δ lead to even greater prevalence of the above correlational patterns).

Since the coefficient of intervention γ_1 is itself a function of $p = \frac{p_X}{p_0 + p_X}$ and δ, it is sufficient to specify values p, v_0 and the variance v_X of X_t in addition to δ to obtain a full parameterization of our model. Finally, Eqs. (6) and (7) indicate that phenotypic correlations are unaffected by a common rescaling of all variances; hence, we can focus on only two parameters, namely p and the ratio of v_0 to v_X. Figure 2 depicts the former on the vertical axis and the latter on the horizontal axis.

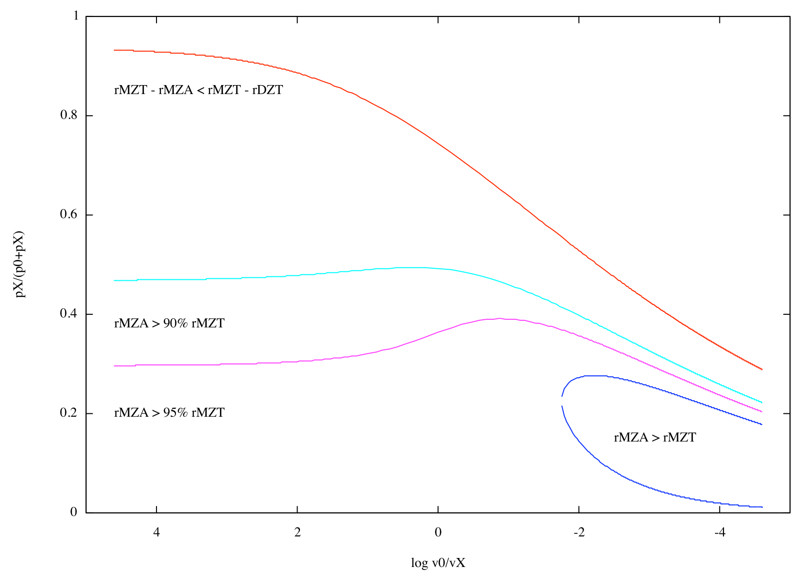

Figure 2: Twins reared together vs. Twins reared apart

Every point below the topmost curve corresponds to a parameterization for which $r^{MZT} - r^{MZA}$ is smaller than $r^{MZT} - r^{DZT}$. The second curve from the top is the upper bound of the region where r^{MZA} is more than 90% of r^{MZT}, and the one immediately below it demarcates the region where r^{MZA} is more than 95% of r^{MZT}. Finally, the oval-shaped region bounded by the last two curves corresponds to parameter values for which r^{MZA} is actually greater than r^{MZT}.

Two main conclusions can be drawn by inspecting Fig. 2. First of all, as noted above, for a substantial range of parameter values, our model generates correlational patterns that the BG literature has interpreted as indicating limited parental influence on developmental outcomes—despite the fact that our model accords a fundamental role to parents.

Second, our setup provides a richer framework than the benchmark ACE model to interpret findings from twin and adoption studies. For instance, suppose that p is greater than $\frac{1}{2}$, indicating a relatively higher weight of experience in the child's learning process—and, therefore, a relatively limited contribution of genetic factors to the phenotype. As Fig. 2 shows, whether, for instance, $r^{MZT} - r^{MZA}$ is smaller or larger than $r^{MZT} - r^{DZT}$ actually depends upon the relative magnitude of the variances v_0 and v_X. If the environment is relatively homogeneous in the population under consideration (i.e. v_X is small relative to v_0), then our model predicts that $r^{MZT} - r^{MZA} < r^{MZT} - r^{DZT}$. Behavioral geneticists would interpret this as indicating that genetic factors are determinant; however, our model suggests that this pattern may instead be entirely due to the relative homogeneity of the environment.

In this respect, our model formalizes an objection to the BG interpretation of correlational findings that other researchers have voiced; for instance, see Ridley [58, pp. 86-87]. Moreover, it also qualifies this objection: if p is small, so that the child's genetically-determined prior has a large weight, then the homogeneity of the environment does not matter: the same correlation patterns will emerge.

3.5 Differences in Phenotypic Variance

The observation that adoptive parents provide more sheltering has a further empirical implication in our model: the ability of adopted children to perform the task under consideration is more dispersed—that is, the variance of the posterior mean is larger for adopted children. This is because adopted children are less exposed to non-shared experiences, and hence don't learn the correct way to perform the task at hand as fast as children reared by their biological parents.

Formally, the analysis in the Appendix shows that the variance of the posterior for a child reared by her biological parents, v_1, and for a child reared by adoptive parents, v_{1a}, satisfy

$$v_1 = v_{1a} - 2\frac{\gamma_1 p_X (p_0 + \gamma_1 p_X)}{(p_0 + p_X)^2} \text{Cov}\left[M_0, Z_0\right],$$

where M_0 and Z_0 refer to a child and her biological parent respectively. The covariance term is multiplied by a negative constant, because $\gamma_1 \in (0, 1)$; this delivers the required conclusion.

However, it is important to point out that this is a statement about the *variance* of posterior means, not their population average. In fact, Eq. (5) implies that the average posterior mean will be the same for adopted and biological children. This is a consequence of the linearity of the parent's optimal policy.

3.6 Heritability Estimates

Within the framework of the ACE model, a key quantity of interest is **heritability**, defined as the ratio

$$\frac{\mathrm{Var}[A]}{\mathrm{Var}[P]} \equiv h^2. \tag{8}$$

Heritability is intended to capture the extent of variation in the observable trait of interest that can be ascribed to variation in the genes. As an example, for several measures of IQ, Bouchard [19] reports values of heritability ranging from 69% to 78%. Other studies have broadly confirmed this finding. According to the above interpretation, these figure suggest that IQ is, to a large extent, genetically determined. It should be noted that the interpretation of heritability adopted in BG is subject to a number of qualifications (see e.g. Goldberger [39]); nevertheless, heritability remains a central quantity of interest in BG, and one that is often invoked in various instances of the nature-nurture debate.

One common measure of heritability is the correlation in the measured trait under consideration for MZ twins raised apart. The intuition suggested by behavioral geneticists is that, due to their being reared in different families and, more broadly, environments, such individuals can only be alike to the extent that their genetic endowment (which is identical) influences the measure of interest.

Another commmon measure of heritability is the difference between the correlations in a measured trait of interest for MZ and DZ twins reared by their biological parents, multiplied by two. The suggested intuition is that taking the difference between these correlations "cancels out" any (additive) effect of common rearing and, more broadly, environmental factors on the trait of interest, thus identifying purely genetic effects. One advantage of this approach relative to the one involving adopted twins is the availability of considerably larger samples. The findings are broadly in line with those reported above.

In the Appendix we verify that the above intuitive considerations can be made rigorous in the ACE model, but require crucial independence assumptions.

In our model, as shown in the Appendix,

$$r^{MZA} = 2(r^{MZT} - r^{DZT}) = \left(\frac{p_0 + \gamma_1 p_X}{p_0 + p_X}\right)^2 h^2.$$

Since $\gamma_1 > 0$, it follows that *standard measures of heritability underestimate the effects of parenting* on the trait of interest.

Slightly rephrasing, standard calculations in BG can be seen as capturing a *broad* notion of heritability, reflecting both the direct impact of genes on the phenotype, as well as their indirect impact, mediated by parenting. However, even comparatively large values of "broad heritability" do not provide any rationale for negating a significant role for parenting as a contributor to developmental outcomes.

4 Concluding Remarks

TO BE REWRITTEN

No direct communication. As noted in §1.1, we do not model direct communication between the parent and the child, because we wish to focus specifically on supervised learning-by-doing.

One way to construct a model that incorporates both aspects of parent/child interaction in a textbook-equilibrium setting is to assume that the parent's objective function does *not* coincide with that of the child, and introduce a round of cheap talk before supervised learning begins. For a simple specification, continue to assume that the child's loss at time t is $(X_t + a_t - b_t)^2$, and suppose that the parent's loss is $(X_t + a_t + D - b_t)^2$. The parameter $D \in \mathbb{R}$ can be interpreted as an *intrinsic* difference between the child's and the parent's preferred way to perform the task at hand, which is distinct from the child's *informational* bias $M_0 - M$.

The analysis and results in this section can be easily adapted to this model; the same is true of our characterization of textbook-equilibrium behavior in Sec. C.3. If we now add one round of cheap talk, the analysis of Crawford and Sobel [30] implies that full information revelation will *not* obtain in equilibrium. Thus, there will be scope for both verbal communication (cheap talk) and supervised learning.

A Appendix

A.1 Main Result

It is convenient to let $p_t = p_0 + tp_X$ for $t = 0, \ldots, L$. Also, to track the child's conditional expectation, define the random variable

$$M_t^a = \begin{cases} \frac{p_0 M_0 + p_X \sum_{s=1}^{t}[X_s + a_s(M, X_1, \ldots, X_{t-1})]}{p_t} & t \leq T \\ \frac{p_0 M_0 + p_X \sum_{s=1}^{T}[X_s + a_s(M, X_1, \ldots, X_{t-1})] + p_X \sum_{s=T+1}^{t} X_s}{p_t} & t > T \end{cases} \tag{9}$$

The formal statement of Theorem 2.1 is as follows:

Theorem A.1 *The optimal parenting policy $a = (a_1, \ldots, a_T) \in \mathcal{A}$ is*

$$a_t = \gamma_t (M_{t-1}^a - M), \tag{10}$$

$$\text{where} \quad \gamma_t = \frac{1 - \delta B_{t+1} \frac{p_X}{p_t} \frac{p_{t-1}}{p_t}}{1 + \delta B_{t+1} \left(\frac{p_X}{p_t}\right)^2}, \quad B_{T+1} = \sum_{\tau=1}^{L-T} \delta^{\tau-1} \left(\frac{p_T}{p_{T+\tau-1}}\right)^2, \quad B_t = \frac{\delta B_{t+1}}{1 + \delta B_{t+1} \left(\frac{p_X}{p_t}\right)^2}. \tag{11}$$

Furthermore, $\gamma_t \in (0,1)$ and γ_t is decreasing in δ and L. Finally, for $t = 1, \ldots, T$,

$$\gamma_t = \mu_t \gamma_t^{\text{FS}} + (1 - \mu_t) \gamma_t^{\text{BC}}, \quad \text{where} \quad \mu_t = \frac{B_t}{\delta B_{t+1}} \in (0,1), \ \gamma_t^{\text{FS}} = 1, \ \gamma_t^{\text{BC}} = -\frac{p_{t-1}}{p_X}.$$

Proof. The fact that γ_t is decreasing in δ and L follows from a simple induction argument, by inspecting Eq. (11).

Henceforth, we let $a = (a_1, \ldots, a_T) \in \mathcal{A}$ denote the parent's optimal policy; also, to simplify the notation, we write M^a_{t-1} simply as M_{t-1}. Thus, by the arguments given in the text, at each time t, the child's optimal action is $b_t = M_{t-1}$.

Begin by analyzing the non-teaching periods. To this end, observe first that, from Eq. (9), for all $t \geq T$ and $\tau \geq 0$,

$$M_{t+\tau} = \frac{1}{p_{t+\tau}} \left(p_t M_t + p_X \sum_{s=t+1}^{t+\tau} X_s \right), \tag{12}$$

where, as is customary, for $\tau = 0$, the empty summation is assumed to equal zero. Hence, for $t \geq T$ and $\tau \geq 1$,

$$\begin{aligned} X_{t+\tau} - M_{t+\tau-1} &= X_{t+\tau} - \frac{p_t}{p_{t+\tau-1}} M_t - \frac{p_X}{p_{t+\tau-1}} \sum_{s=t+1}^{t+\tau-1} X_s = \\ &= (X_{t+\tau} - M) - \frac{p_t}{p_{t+\tau-1}} (M_t - M) - \frac{p_X}{p_{t+\tau-1}} \sum_{s=t+1}^{t+\tau-1} (X_s - M); \end{aligned} \tag{13}$$

the last line uses the fact that $p_t M + p_X \sum_{s=t+1}^{t+\tau-1} M = p_t M + p_X(\tau - 1)M = p_{t+\tau-1}M$. It now follows that, at any time $t \geq T$, and for all $\tau \geq 1$, from the point of view of the parent, i.e. conditional upon the realization of M, the expected loss at time $t + \tau$ given the observed value of M_t is

$$\mathrm{E}[(X_{t+\tau} - M_{t+\tau-1})^2 | M, M_t] = \frac{1}{p_X} + \frac{p_X^2}{p_{t+\tau-1}^2} (\tau - 1) \frac{1}{p_X} + \frac{p_t^2}{p_{t+\tau-1}^2} (M_t - M)^2;$$

this follows from the assumption that $X_1, \ldots, X_L$ are i.i.d. normal with mean M and precision p_X conditional on M,[27] which in turn implies that all of the cross-terms, which are of the form $(X_{t+\tau} - M)(M_t - M)$, $(X_{t+\tau} - M)(X_s - M)$ and $(M_t - M)(X_s - M)$ for $s \in \{t+1, \ldots, t+\tau-1\}$, and $(X_s - M)(X_\sigma - M)$ for s, σ distinct in $\{t + 1, \ldots, t + \tau - 1\}$, all have zero conditional expectation.

It follows that, for every non-teaching period $t + 1 \in \{T + 1, \ldots, L\}$, conditional upon M and M_t, the expected time-$(t + 1)$ continuation value of the child's optimal policy is $V_{t+1}(M_t, M) = A_{t+1} + B_{t+1}(M_t - M)^2$, where[28]

$$A_{t+1} = \sum_{\tau=1}^{L-t} \delta^{\tau-1} \left[\frac{1}{p_X} + \frac{p_X}{p_{t+\tau-1}^2} (\tau - 1) \right], \quad B_{t+1} = \sum_{\tau=1}^{L-t} \delta^{\tau-1} \left(\frac{p_t}{p_{t+\tau-1}} \right)^2. \tag{14}$$

Turn now to teaching periods $t \in \{1, \ldots, T\}$. From the argument just given, $V_{T+1}(M_T, M) = A_T + B_{T+1}(M_T - M)^2$. We now show inductively that, for $t = T, \ldots, 1$, if $V_{t+1}(M_t, M) = A_{t+1} +$

[27]Strictly speaking, since the parent knows M, in her view this is the actual unconditional distribution of the X_s's; for the child, M is a r.v., so the above statement describes the conditional distribution of the X_s's given M. However, we use a more explicit terminology and notation to remind the reader that M can be treated as a constant in the above expectations.

[28]For the purposes of solving for the optimal policy, we only need an expression for V_{T+1}. However, the expression for B_t, $t = T + 1, \ldots, L$ is used below to show that $\gamma_t \in (0, 1)$ in the scalar case.

$B_{t+1}(M_t - M)^2$, then the equations for a_t and γ_t in Thm. 2.1 hold, and furthermore $V_t(M_{t-1}, M) = A_t + B_t(M_{t-1} - M)^2$, where B_t is again as in Thm. 2.1.

Notice that, again from Eq. (9), for every $t \leq T$,

$$M_t = \frac{p_0}{p_t} M_0 + \frac{p_X}{p_t} \sum_{s=1}^{t} (X_s + a_s) = \frac{p_{t-1}}{p_t} M_{t-1} + \frac{p_X}{p_t}(X_t + a_t). \tag{15}$$

After substituting for M_t in the expression for V_{t+1} in the inductive hypothesis, conditional on the information $\mathcal{I}_t \equiv \{M, X_1, \ldots, X_{t-1}\}$, the action a_t must solve the Bellman equation

$$\begin{aligned} V_t(M_{t-1}, M) &= \min_{\bar{a}} \mathrm{E}\left[(X_t + \bar{a} - M_{t-1})^2 | \mathcal{I}_t\right] + \\ &+ \delta \mathrm{E}\left[A_{t+1} + B_{t+1}\left(\frac{p_{t-1}}{p_t} M_{t-1} + \frac{p_X}{p_t}(X_t + \bar{a}) - M\right)^2 | \mathcal{I}_t\right]. \end{aligned} \tag{16}$$

Differentiating with respect to $\bar{a}$, taking expectations, and dividing by 2 yields the FOC

$$0 = \bar{a} - (M_{t-1} - M) + \delta B_{t+1}\left(\frac{p_{t-1}}{p_t} M_{t-1} + \frac{p_X}{p_t}(M + a) - M\right)\frac{p_X}{p_t}$$

and therefore

$$a_t = \frac{1 - \delta B_{t+1} \frac{p_X p_{t-1}}{p_t^2}}{1 + \delta B_{t+1}\left(\frac{p_X}{p_t}\right)^2} \cdot (M_{t-1} - M) \equiv \gamma_t (M_{t-1} - M). \tag{17}$$

We now show that V_t can be expressed as a quadratic form in $(M_{t-1} - M)$. First, to simplify the quadratic form in the first line of Eq. (16), observe that

$$X_t + \gamma_t(M_{t-1} - M) - M_{t-1} = (X_t - M) - (1 - \gamma_t)(M_{t-1} - M);$$

as for the quadratic form in the second line of Eq. (16),

$$\begin{aligned} &\frac{p_{t-1}}{p_t} M_{t-1} + \frac{p_X}{p_t}(X_t + \gamma_t(M_{t-1} - M)) - M = \\ &= \frac{p_{t-1}}{p_t}(M_{t-1} - M) + \frac{p_X}{p_t}(X_t - M + \gamma_t(M_{t-1} - M)) = \\ &= \frac{p_{t-1} + p_X \gamma_t}{p_t}(M_{t-1} - M) + \frac{p_X}{p_t}(X_t - M). \end{aligned}$$

Therefore $V_t(M_{t-1}, M) = A_t + B_t(M_{t-1} - M)^2$, where A_t is a suitable constant and

$$B_t = (1 - \gamma_t)^2 + \delta B_{t+1}\left(\frac{p_{t-1} + p_X \gamma_t}{p_t}\right)^2. \tag{18}$$

To show that B_t can be written as in Eq. (11), note that

$$1 - \gamma_t = \frac{1 + \delta B_{t+1}\left(\frac{p_X}{p_t}\right)^2 - 1 + \delta B_{t+1}\frac{p_X p_{t-1}}{p_t^2}}{1 + \delta B_{t+1}\left(\frac{p_X}{p_t}\right)^2} = \frac{\delta B_{t+1}}{1 + \delta B_{t+1}\left(\frac{p_X}{p_t}\right)^2}\frac{p_X}{p_t} \tag{19}$$

and, similarly,

$$\frac{p_{t-1}+\gamma_t p_X}{p_t} = \frac{1}{p_t}\frac{p_{t-1}+p_{t-1}\delta B_{t+1}\left(\frac{p_X}{p_t}\right)^2+p_X-p_X\delta B_{t+1}\frac{p_X p_{t-1}}{p_t^2}}{1+\delta B_{t+1}\left(\frac{p_X}{p_t}\right)^2} = \frac{1}{1+\delta B_{t+1}\left(\frac{p_X}{p_t}\right)^2}. \tag{20}$$

Hence, as required,

$$B_t = \frac{(\delta B_{t+1})^2\left(\frac{p_X}{p_t}\right)^2+\delta B_{t+1}}{\left[1+\delta B_{t+1}\left(\frac{p_X}{p_t}\right)^2\right]^2} = \frac{\delta B_{t+1}}{1+\delta B_{t+1}\left(\frac{p_X}{p_t}\right)^2}.$$

We now show that $\gamma_t \in (0,1)$ for $t = 1,\dots,T$. It is clear that $\gamma_t < 1$, so we must only verify that $\gamma_t > 0$. Notice first that, from Eq. (14), for $t = T+1,\dots,L-1$,

$$\begin{aligned} B_t &= \sum_{\tau=1}^{L-(t-1)}\delta^{\tau-1}\left(\frac{p_{t-1}}{p_{t+\tau-2}}\right)^2 = \left(\frac{p_{t-1}}{p_{t-1}}\right)^2+\sum_{\tau=2}^{L-(t-1)}\delta^{\tau-1}\left(\frac{p_{t-1}}{p_{t+\tau-2}}\right)^2 = \\ &= 1+\delta\sum_{\tau=1}^{L-t}\delta^{\tau-1}\left(\frac{p_{t-1}}{p_{t+\tau-1}}\right)^2 = 1+\delta\left(\frac{p_{t-1}}{p_t}\right)^2\sum_{\tau=1}^{L-t}\delta^{\tau-1}\left(\frac{p_t}{p_{t+\tau-1}}\right)^2 = \\ &= 1+\delta\left(\frac{p_{t-1}}{p_t}\right)^2 B_{t+1}; \end{aligned}$$

also, $B_L = 1$. Furthermore, we claim that $B_{t+1}\frac{p_X}{p_t}\frac{p_{t-1}}{p_t} < 1$ for all $t < L$. The claim is true for $t = L-1$, because $B_L = 1$ and $p_t = p_X + p_{t-1}$ for all $t \geq 1$. Now consider an arbitrary $t \in \{T+1,\dots,L\}$ and assume the claim is true for $t+1,\dots,L$. Now $B_t\frac{p_X}{p_{t-1}}\frac{p_{t-2}}{p_{t-1}} = \frac{p_X}{p_{t-1}}\frac{p_{t-2}}{p_{t-1}} + \delta B_{t+1}\left(\frac{p_{t-1}}{p_t}\right)^2\frac{p_X}{p_{t-1}}\frac{p_{t-2}}{p_{t-1}} = \frac{p_X}{p_{t-1}}\frac{p_{t-2}}{p_{t-1}} + \delta B_{t+1}\frac{p_X}{p_t}\frac{p_{t-1}}{p_t}\cdot\frac{p_{t-1}}{p_{t-1}}\frac{p_{t-2}}{p_{t-1}} \leq \frac{p_X}{p_{t-1}} + \delta B_{t+1}\frac{p_X}{p_t}\frac{p_{t-1}}{p_t}\cdot\frac{p_{t-2}}{p_{t-1}} < 1$.

Hence, in particular, $\delta B_{T+1}\frac{p_X}{p_T}\frac{p_{T-1}}{p_T} < 1$, and so $\gamma_T \in (0,1)$. If $T = 1$, we are done. Otherwise, argue by induction; consider $t \in \{1,\dots,T\}$ and assume that the claim is true for $t+1$. By Eqs. (19) and (11), $B_{t+1} = (1-\gamma_{t+1})\frac{p_{t+1}}{p_X}$, so $B_{t+1}\frac{p_X}{p_t}\frac{p_{t-1}}{p_t} = (1-\gamma_{t+1})\frac{p_{t+1}}{p_X}\frac{p_X}{p_t}\frac{p_{t-1}}{p_t} = (1-\gamma_{t+1})\frac{(p_t+p_X)(p_t-p_X)}{p_t^2} = (1-\gamma_{t+1})\frac{p_t^2-p_X^2}{p_t^2} = (1-\gamma_{t+1})\left(1-\frac{p_X^2}{p_t^2}\right) < (1-\gamma_{t+1})$, because $p_t = p_0 + tp_X \geq p_0 + p_X$ for $t \geq 1$. By the induction hypothesis, $1-\gamma_{t+1} \in (0,1)$, so $B_{t+1}\frac{p_X}{p_t}\frac{p_{t-1}}{p_t} < 1$ and therefore $\gamma_t \in (0,1)$, as claimed.

Finally, we show that γ_t can be decomposed as a weighted average of γ_t^{FS} and γ_t^{BC}. We have

$$\begin{aligned} \gamma_t &= \frac{1-\delta B_{t+1}\frac{p_X}{p_t}\frac{p_{t-1}}{p_t}}{1+\delta B_{t+1}\left(\frac{p_X}{p_t}\right)^2} = \frac{1+\delta B_{t+1}\left(\frac{p_X}{p_t}\right)^2\left(-\frac{p_{t-1}}{p_X}\right)}{1+\delta B_{t+1}\left(\frac{p_X}{p_t}\right)^2} = \\ &= \frac{1}{1+\delta B_{t+1}\left(\frac{p_X}{p_t}\right)^2}\cdot 1 + \frac{\delta B_{t+1}\left(\frac{p_X}{p_t}\right)^2}{1+\delta B_{t+1}\left(\frac{p_X}{p_t}\right)^2}\cdot\left(-\frac{p_{t-1}}{p_X}\right) = \\ &\equiv \mu_t\gamma_t^{\mathrm{FS}} + (1-\mu_t)\gamma_t^{\mathrm{BC}}, \end{aligned}$$

and Eq. (11) implies that

$$\mu_t = \frac{1}{1 + \delta B_{t+1} \left(\frac{p_X}{p_t}\right)^2} = \frac{B_t}{\delta B_{t+1}}.$$

which clearly lies in $(0, 1)$. ■

A.2 Time pattern of μ_t

Proposition A.2 *In the setting of Theorem 2.1, for every $T_0 \in \{1, \ldots, T\}$:*

1. $\delta < \left(\frac{p_{T_0}}{p_{T_0+1}}\right)^2$ *implies $\mu_t > \mu_{t+1}$ for all $t = T_0, \ldots, T-1$.*
2. *There exists $\delta(T_0) \in (0, 1)$ such that $\delta > \delta(T_0)$ implies $\mu_t < \mu_{t+1}$ for all $t = T_0, \ldots, T-1$; furthermore, $\delta(t) \le \delta(t+1)$ for all $t = 1, \ldots, T-1$.*

Proof. For the first claim, recall that $\mu_t = \frac{1}{1+\delta B_{t+1}\left(\frac{p_X}{p_t}\right)^2}$ and

$$B_{t+1}\left(\frac{p_X}{p_t}\right)^2 = \left(\frac{p_{t+1}}{p_t}\right)^2 \frac{\delta}{1 + \delta B_{t+2}\left(\frac{p_X}{p_{t+1}}\right)^2} B_{t+2}\left(\frac{p_X}{p_{t+1}}\right)^2;$$

Clearly, if $\delta < \left(\frac{p_{T_0}}{p_{T_0+1}}\right)^2$, then $B_{T_0+1}\left(\frac{p_X}{p_{T_0}}\right)^2 < B_{T_0+2}\left(\frac{p_X}{p_{T_0+1}}\right)^2$ and so $\mu_{T_0} < \mu_{T_0+1}$; since $\frac{p_t}{p_{t+1}}$ is increasing in t, we have $\mu_t > \mu_{t+1}$ for all $t = T_0, \ldots, T$, as required.

For the second claim, recall from Eq. (18) that

$$B_{t+1} = (1 - \gamma_{t+1})^2 + \delta B_{t+2}\left(\frac{p_t + \gamma_{t+1} p_X}{p_{t+1}}\right)^2.$$

Multiplying both sides by $\left(\frac{p_X}{p_t}\right)^2$ and expanding the square in the second term in the rhs yields

$$B_{t+1}\left(\frac{p_X}{p_t}\right)^2 = (1-\gamma_{t+1})^2\left(\frac{p_X}{p_t}\right)^2 + \delta B_{t+2}\left(\frac{p_X}{p_{t+1}}\right)^2 + \delta B_{t+2}\left(\frac{p_X}{p_t}\right)^2\left[\left(\frac{\gamma_{t+1} p_X}{p_{t+1}}\right)^2 + \left(\frac{2 p_t \gamma_{t+1} p_X}{p_{t+1}}\right)\right].$$

Now observe that, by Eq. (11), $\gamma_{t+1} \in (0, 1)$ also for $\delta = 1$. Hence, the above equation shows that, for $\delta = 1$, $B_{t+1}\left(\frac{p_X}{p_t}\right)^2 > B_{t+2}\left(\frac{p_X}{p_{t+1}}\right)^2$, and hence also $\mu_t < \mu_{t+1}$. Since both sides of the latter equation are continuous in δ, there is $\delta_t \in [0, 1)$ such that $\delta > \delta_t$ implies $\mu_t < \mu_{t+1}$. Now take $\delta(T_0) = \min\{\delta_{T_0}, \delta_{T_0+1}, \ldots, \delta_T\}$ to obtain a cutoff with the required properties. ■

A.3 Correlations and Heritability in the ACE model

This subsection is based on Plomin et al. [55] (see also Goldberger [39]).

Notation: the variables for members of a twin pair are denoted by P, A, C, E and P', A', C', E' respectively. If X and Y denote any of these variables, denote by v_X the variance of X and by r_{XY} the correlation between X and Y.

If genetype, shared envoronment, and non-shared environment are mutually independent and identically distributed across twins, so that $v_A = v_{A'}$, etc., then clearly

$$v_P = v_{P'} = v_A + v_C + v_E = v_{A'} + v_{C'} + v_{E'},$$

regardless of zygosity or rearing conditions. Furthermore, in general,

$$r_{PP'} = r_{AA'}\frac{v_A}{v_P} + r_{CC'}\frac{v_C}{v_P} + r_{EE'}\frac{v_E}{v_P}.$$

Now, by definition, the non-shared environmental factors are uncorrelated across twins, so $r_{EE'} = 0$. MZ twins share the same genetic endowment, so $A = A'$ and $r_{AA'} = 1$; DZ twins share only half of the genes, so $r_{AA'} = \frac{1}{2}$. Finally, twins reared together by definition share the same "common environment," so $r_{CC'} = 1$ for these twins; on the other hand, for twins reared apart, $r_{CC'} = 0$. Recalling that $h^2 = \frac{v_A}{v_P}$ and similarly defining $c^2 = \frac{v_C}{v_P}$, we have

$$r^{MZT} = h^2 + c^2, \quad r^{MZA} = h^2, \quad r^{DZT} = \frac{1}{2}h^2 + c^2.$$

It is now immediate to see that, as we claim in the text:

1. $r^{MZT} - r^{MZA} = c^2$ captures the effect of common rearing on phenotypic correlation; notice that this quantity must necessarily be positive;
2. $r^{MZT} - r^{MZA} < r^{MZT} - r^{DZT}$ iff $c^2 < \frac{1}{2}h^2$, i.e. this inequality indicates that genes matter more than twice as much as commor rearing; and finally
3. $2\left[r^{MZT} - r^{DZT}\right] = r^{MZA} = h^2$ yields an estimate of heritability.

A.4 Supervised Learning and Twin Correlations

Throughout this subsection, we continue to let $p = \frac{p_X}{p_0 + p_X}$ as in §3.4. We begin by calculating phenotypic variance for children reared by their biological parents: from Eq. (5), since X_1 is uncorrelated with M_0 and Z_0, we get

$$v_1 = [(1-p) + p\gamma_1]^2 v_0 + p^2 v_X + [p\gamma_1]^2 v_{Z0} - 2p(1-p)r_0\sqrt{v_0 v_{Z0}}. \tag{21}$$

For children reared by adoptive parents, the correlation of M_0 and Z_0 is zero, so

$$v_{1a} = [(1-p) + p\gamma_1]^2 v_0 + p^2 v_X + [p\gamma_1]^2 v_{Z0}. \tag{22}$$

We now turn to correlations. Let M_1 and M_1' be the posterior of two twins. For MZ twins reared together, Eq. (5) yields $M_1 = [(1-p)+p\gamma_1]M_0 - p\gamma_1 Z_0 + pX_1$ and $M_1' = [(1-p)+p\gamma_1]M_0 - p\gamma_1 Z_0 + pX_1'$, reflecting the fact that these individuals have the same genetic endowment and the same parents. It follows that

$$r^{MZT} = [(1-p) + p\gamma_1]^2\frac{v_0}{v_1} + [p\gamma_1]^2\frac{v_{Z0}}{v_1} - 2[(1-p) + p\gamma_1]p\gamma_1 r_0\frac{\sqrt{v_0 v_{Z0}}}{v_1}. \tag{23}$$

For DZ twins reared together, we have $M_1 = [(1-p)+p\gamma_1]M_0 - p\gamma_1 Z_0 + pX_1$ and $M_1' = [(1-p)+p\gamma_1]M_0' - p\gamma_1 Z_0 + pX_1'$, reflecting the fact that these individuals share the same parents; their genetic endowment is now different but biological considerations suggest that the correlation between M_0 and M_0' is $\frac{1}{2}$. Therefore,

$$r^{DZT} = \frac{1}{2}[(1-p)+p\gamma_1]^2 \frac{v_0}{v_1} + [p\gamma_1]^2 \frac{v_{Z0}}{v_1} - 2[(1-p)+p\gamma_1]p\gamma_1 r_0 \frac{\sqrt{v_0 v_{Z0}}}{v_1}, \tag{24}$$

which, together with Eq. (23, immediately yields Eq. (7) in the text. Finally, for MZ twins reared apart, we have $M_1 = [(1-p)+p\gamma_1]M_0 - p\gamma_1 Z_0 + pX_1$ and $M_1' = [(1-p)+p\gamma_1]M_0 - p\gamma_1 Z_0' + pX_1'$, reflecting the fact that these individuals have the same genetic endowment, but different parents. Furthermore, if parents are independently drawn from the population of interest, the correlation of Z_0 and Z_0' is zero, and so

$$r^{MZA} = [(1-p)+p\gamma_1]^2 \frac{v_0}{v_{1a}}. \tag{25}$$

We now have

$$\begin{aligned}
r^{MZT} - r^{MZA} &= [(1-p)+p\gamma_1]^2 \left(\frac{v_0}{v_1} - \frac{v_0}{v_{1a}}\right) + [p\gamma_1]^2 \frac{v_{Z0}}{v_1} - 2[(1-p)+p\gamma_1]p\gamma_1 r_0 \frac{\sqrt{v_0 v_{Z0}}}{v_1} = \\
&= [(1-p)+p\gamma_1]^2 v_0 \frac{v_{1a}-v_1}{v_1 v_{1a}} + [p\gamma_1]^2 \frac{v_{Z0}}{v_1} - 2[(1-p)+p\gamma_1]p\gamma_1 r_0 \frac{\sqrt{v_0 v_{Z0}}}{v_1} = \\
&= -\left\{1 - [(1-p)+p\gamma_1]^2 \frac{v_0}{v_{1a}}\right\} 2[(1-p)+p\gamma_1]p\gamma_1 r_0 \frac{\sqrt{v_0 v_{Z0}}}{v_1} + [p\gamma_1]^2 \frac{v_{Z0}}{v_1} = \\
&= -2\frac{v_X + \gamma_1^2 v_{Z0}}{v_{1a}}[(1-p)+p\gamma_1]p^3\gamma_1 r_0 \frac{\sqrt{v_0 v_{Z0}}}{v_1} + [p\gamma_1]^2 \frac{v_{Z0}}{v_1}
\end{aligned}$$

which is Eq. (6); the penultimate step follows from $v_{1a} - v_1 = 2[(1-p)+p\gamma_1]p\gamma_1 r_0 \sqrt{v_0 v_{Z0}}$, and the last step from Eq. (22).

Finally, we describe the steady-state model we used to generate Fig. 2. We shall first compute Z_0 as the posterior of an individual with prior M_{-1}, who learns from the observation of X_{-1} under the supervision of a *biological* parent[29] with time-(-1) mean Z_{-1}. By analogy with Eq. (5),

$$Z_0 = [(1-p)+p\gamma_1]M_{-1} + pX_0 - p\gamma_1 Z_{-1};$$

we can now write Z_{-1} similarly as

$$Z_{-1} = [(1-p)+p\gamma_1]M_{-2} + pX_{-1} - p\gamma_1 Z_{-2},$$

and use this expression to substitute Z_{-1} above:

$$Z_0 = [(1-p)+p\gamma_1]M_{-1} + pX_0 - p\gamma_1[(1-p)+p\gamma_1]M_{-2} - p^2\gamma_1 X_{-1} + [p\gamma_1]^2.Z_{-2}$$

Thus, by iterating this process,

$$Z_0 = \sum_{t=1}^{\infty} (-p\gamma_1)^{t-1} \left\{[(1-p)+p\gamma_1]M_{-t} + pX_{-(t-1)}\right\}.$$

[29] Assuming that a small fraction of children are adopted in each generation does not change our findings significantly.

Now a child shares approximately 50% of her genes with her parent, who in turn shares approximately 50% of his genes with his own parent, etc.; thus, the correlation between M_0 and M_{-t} can be taken to be 2^{-t}, so their covariance is $2^{-t}v_0$. Furthermore, M_0 is uncorrelated with observations; thus,

$$\mathrm{Cov}\,[M_0, Z_0] = [(1-p) + p\gamma_1]v_0 \sum_{t=1}^{\infty}(-p\gamma_1)^{t-1}\left(\frac{1}{2}\right)^t.$$

Finally, to compute v_{Z0}, note that, in steady state, it must be the case that $v_1 = v_{Z0}$; furthermore, clearly $\mathrm{Cov}\,[M_0, Z_0] = r_0\sqrt{v_0 v_{Z0}}$. Thus, we can substitute for the latter quantity in Eq. (21), assume that $v_1 = v_{Z0}$ and solve for v_{Z0}: we get

$$v_{Z0} = v_1 = \frac{[(1-p) + p\gamma_1]^2 v_0 + p^2 v_X - 2p(1-p)\mathrm{Cov}\,[M_0, Z_0]}{1 - [p\gamma_1]^2}.$$

References

[1] Apesteguia, Jose, Steffen, and Oechssler, Jrg(2006): "Imitation: theory and experimental evidence" *Journal of Economic Theory*, forthcoming.

[2] Altmann (1980) *Baboon Mothers and Infants* The University of Chicago Press.

[3] Bandura A. (1977): *Social Learning Theory*, Prentice-Hall.

[4] Bandura A. and R.H. Walters (1963): *Social Learning and Personality Development*, New York: Holt, Rinehart & Winston.

[5] Baumrind, D. (1967): "Child care practices anteceding three patterns of preschool behavior," *Genetic Psychology Monographs*, 75, 43-88.

[6] Baydar, N. and Brooks-Gunn, J. (1991): "Effects of maternal employment and childcare arrangements on preschoolers cognitive and behavioral outcomes: evidence from the Children of the NLSY." *Developmental Psychology*, vol. 27, 932-45.

[7] Becker, G. S. (1981): *A Treatise on the Family*. Harvard University Press, Cambridge, Massachusetts.

[8] Belsky, J. and Eggebeen, D. (1991): "Early and extensive maternal employment and young children's socioemotional development: Children of the National Longitudinal Survey of Youth." *Journal of Marriage and the Family*, vol. 53, 1083-1110.

[9] Benabou, Roland and Jean Tirole (2002) "Self-Confidence and Personal Motivation" *Quarterly Journal of Economics,* 117(3), 871-915.

[10] Benabou, Roland and Jean Tirole (2003) "Intrinsic and Extrinsic Motivation" *Review of Economic Studies,* 70(3) (2003), 489-520.

[11] Berardo, J., and Smith, A. (2000): *Bayesian Theory*, 2nd ed., John Wiley and Sons, New York.

[12] Bernal, Raquel (2004): "Employment and Child Care Decisions of Mother's and the Well-being of their Children," manuscript, Northwestern University.

[13] Bisin, A. and T. Verdier (2000): "Beyond the Melting Pot: Cultural Transmission, Marriage, and the Evolution of Ethnic and Religious Traits," *Quarterly Journal of Economics*, CXV(3), 955-988.

[14] Bisin, A. and T. Verdier (2001): "Cultural Transmission, Socialization, and the Dynamics of Preferences," *Journal of Economic Theory*, 97(1), 298-319.

[15] Bisin, A., G. Topa and T. Verdier (2004): "Religious Intermarriage and Socialization in the U.S.", *Journal of Political Economy*, 112-3, 615-64.

[16] Bjorklund, Anders, Lindahl, Michael and Erik Plug (2006): "The Origins of Intergenerational Associations: Lessons from Swedish Adoption Data," *Quarterly Journal of Economics*, vol. CXXI n. 3, pp. 999-1029.

[17] Black, Sandra, Devereux, Paul, and Salvanes, Kjell (2005): "The More the Merrier? The Effect of Family Size and Birth Order on Childrens Education," *Quarterly Journal of Economics*, Vol. 120, N. 2 669-700.

[18] Blurton-Jones, Nicholas: "The lives of hunter - gatherer children" in *Juvenile primates : life history, development, and behavior* edited by Michael E. Pereira, Lynn A. Fairbanks; New York : Oxford University Press, 1993.

[19] Bouchard, T.J. Jr., Lykken, D.T., McGue, M., Segal, N.L., and Tellegen, A. (1990): "Sources of Human Psychological Differences: The Minnesota Study of Twins Reared apart," *Science*, 250, 223-338.

[20] Boyd, R. and P. Richerson (1985): *Culture and the Evolutionary Process*, Chicago, University of Chicago Press.

[21] Brezzi, M. and Lai, T.L. (2000): "Incomplete Learning from Endogenous Data in Dynamic Allocation," *Econometrica* vol. 68 n. 6, 1511-1516.

[22] Brunnermeier, Markus and Jonathan Parker (2004). "Optimal Expectations," Manuscript, Princeton University.

[23] Camerer, Colin F. (2003): *Behavioral Game Theory: Experiments in Strategic Interaction*; Princeton University Press.

[24] Carneiro, Pedro, Heckman, James J. and Dimitriy V. Masterov (2002): "Labor Market Discrimination and Racial Differences in Premarket Factors," manuscript, University of Chicago.

[25] Cavalli-Sforza, L. and M. Feldman (1981): *Cultural Transmission and Evolution: A Quantitative Approach*, Princeton NJ, Princeton University Press.

[26] Chisholm, K. "A Three Year Follow-up of Attachment and Indiscriminate Friendliness in Children Adopted from Romanian Families." *Child Developmen* 69(4): 1092-1106.

[27] Clutton-Brock, T. H. (1991): *The evolution of parental care* Princeton, N.J. : Princeton University Press.

[28] Collins, W.A., Maccoby, E., Steinberg, L., Hetherington, E.M., and Bornstein, M.H. (2000): "Contemporary research on parenting: the case for nature and nurture," *American Psychologist*, **55** (2), 218-232.

[29] Compte, Olivier and Andrew Postlewaite: "Confidence-Enhanced Performance," *American Economic Review*, forthcoming.

[30] Crawford, Vincent and Joel Sobel (1982): "Strategic Information Transmission," *Econometrica*, 50 n. 6, pp. 1431-1451.

[31] Cunha, Flavio and James Heckman (2007): "The Technology of Skill Formation," NBER Working Paper 12840.

[32] Demo, D. H., and Cox, M. J. (2000): "Families with young children: A review of research in the 1990s," *Journal of Marriage and the Family*, 62, 867-895.

[33] Desai, S., Chase-Lansdale, P. L. and Michael, R. T. (1989): "Mother or market? effects of maternal employment on the intellectual ability of 4-year old children." *Demography*, vol. 26, 545-61.

[34] Easley, David and Nicholas Kiefer "Controlling a Stochastic Process with Unknown Parameters;" *Econometrica*, 1988.

[35] Ettinger, David and Philippe Jehiel (2006): "Towards a Theory of Deception," manuscript, PSE and University College London.

[36] Fischoff, B., Slovic, P., and Lichtenstein, S. (1977): "Knowing with Certainty: The Appropriateness of Extreme Confidence," *Journal of Experimental Psychology: Human Perception and Performance*, 3(4), 552-564.

[37] Fudenberg Drew and David K. Levine (1998): *Theory of Learning in Games*; MIT Press.

[38] Gilboa, I., and D. Schmeidler, "Case-Based Decision Theory", *The Quarterly Journal of Economics*, **110** (1995) 605-639

[39] Goldberger, Arthur S. (1979), "Heritability," *Economica*, New Series, Vol. 46 n. 184, pp. 327–347.

[40] Gopnick, A., and Meltzoff, A. (1992): "Categorization and naming: basic-level sorting in eighteen-month-olds and its relationship to language," *Child Development*, **63**, 1091-1103.

[41] Griffin, Gary A. and Harry F. Harlow (1966) "Effects of Three Months of Total Social Deprivation on Social Adjustment and Learning in the Rhesus Monkey," *Child Development,* Vol. 37, No. 3. pp. 533-547.

[42] Harlow, H. F., and Zimmerman, R. (1959). "Affectional responses in the infant monkey." *Science,* 130, 421-432.

[43] Harbaugh, W., K. Krause, and T. Berry (2001) "AGARP for Kids: On the Development of Rational Choice Behavior," *American Economic Review*, 91 #5, Dec., 1539-1545.

[44] Harris, Judith Rich (1995): "Where Is the Child's Environment? A Group Socialization Theory of Development," *Psychological Review*, Vol. 102, N. 3, 458-489.

[45] Harris, Judith Rich (1998): *The Nurture Assumption.* The Free Press.

[46] Hoopes, Janet L. (1982): *Prediction in Child Development: A Longitudinal Study of Adoptive and Non-Adoptive Families. The Delaware Family Study.* Child Welfare Leagues of America, New York, NY.

[47] Harsanyi, J. (1967-68): "Games of Incomplete Information Played by Bayesian Players. Parts I, II, III," *Management Science*, 14, 159-182, 320-334, 486-502.

[48] Kalai, E., and Lehrer, E. (1993): "Rational learning leads to Nash equilibrium," *Econometrica* **61** (5), 1019-1045.

[49] Koszegi, Botond (2000) "Ego Utility, Overconfidence, and Task Choice," manuscript, University of California, Berkeley.

[50] Lamport-Commons, M., and Miller, P. (1992): "Emotional learning in infants: a cross-cultural examination,"*Developmental Psychology*, **28** n.4.

[51] Lancaster Jane B., and Chet S. Lancaster (1987): "The Watershed: Change in Parental-Investment and Family Formation Strategies in the Course of Human Evolution." in *Parenting Across the Lifespan* edited by Jane B. Lancaster, Jeanne Altmann, Alice S. Rossi, And Lonnie Sherrod; Aldyne de Gruyter, Hawthorne, NY.

[52] OECD Publishing (2006): "Education at a Glance: OECD Indicators - 2005 Edition."

[53] Pinker, S. (2002): *The Blank Slate: The Modern Denial of Human Nature*, Viking-Penguin.

[54] Plomin, R. and Daniels, D. (1987): "Why are children in the same family so different from each other?" *The Behavioral and Brain Sciences,* 10, 1-16.

[55] Plomin, R., DeFries, J. C., McClearn, G. E., & McGuffin, P. (2001): *Behavioral Genetics.* (4th Edition). New York: Worth Publishers.

[56] Reiss David, Neiderheiser Jenae M., Hetherington Mavis E. and Plomin Robert (2000): *The Relationship Code.* Harvard University Press.

[57] Richman, A.L., LeVine, R., Stables, R., Howrigan, G., Welles-Nystrom, B., and LeVine, S. (1988): "Maternal behavior to infants in five cultures,"in *Parental Behavior in Diverse Societies*, LeVine, R., Miller, P., and West, M. (eds.), Jossey-Bass Inc., San Francisco

[58] Ridley, Matt, (2003): *Nature Via Nurture : Genes, Experience, and What Makes Us Human*, Harper Collins.

[59] Robson, Arthur, Kaplan, Hillard, Lancaster, Jane (2006): "Embodied Capital and the Evolutionary Economics Of the Human Lifespan," *Population and Development Review*, forthcoming.

[60] Rogoff, B. (1997): "Cognition as a collaborative process," in *Handbook of Child Psychology, Vol. 2*, Damon, W. (ed.), fifth edition, John Wiley and Sons, New York.

[61] Rubinstein, Ariel (1989) "Optimal Flexibility of Rules: The Tale of the Wise Principal and the Naive Agent," *European Journal of Political Economy* 5 219-227.

[62] Sacerdote, Bruce (2002): "The Nature and Nurture of Economic Outcomes," *American Economic Review Papers and Proceedings*, vol. XCII p. 344–348.

[63] Sacerdote, Bruce (2006): "How Large are the Effects from Changes in the Family Environment? A Study of Korean-American Adoptees," *Quarterly Journal of Economics*, forthcoming.

[64] Shonkoff, Jack, and Deborah Phillips, eds. 2000. *From Neurons to Neighborhoods: The Science of Early Childhood Development*. Washington, D.C.: National Academy Press.

[65] Smith, Lones and Peter Sorensen (2006): "Informational Herding and Optimal Experimentation," manuscript, University of Michigan.

[66] Schlag, Karl H. (1998) "Why Imitate, and If So, How?, : A Boundedly Rational Approach to Multi-armed Bandits" *Journal of Economic Theory*, Volume 78, Issue 1, Pages 130-156

[67] Todd, Petra E., Kenneth I. Wolpin, (2004). "The Production of Cognitive Achievement in Children: Home, School and Racial Test Score Gaps" manuscript, University of Pennsylvania.

[68] Turkheimer, Eric, Haley, Andreanna, Waldron, Mary, D'Onofrio, Brian, and Gottesman, Irving I. (2003): "Socioeconomic Status Modifies Heritability of IQ in Young Children," *Psycological Science* XIV 623-628

[69] Vandell D.L. and J. Ramanan (1992) "Effects of Early and Recent Maternal Employment on Children from Low Income Families." *Child Development* 63: 938-949.

[70] Vygotsky L.S. (1978). *Mind in Society.* Cambridge, MA: Harvard University Press.

[71] Waxman, S. (in press): "Links between object categorization and naming: Origins and emergence in human infants," In D. K. Rakison & L. M. Oakes (Eds.), *Early category and concept development: Making sense of the blooming, buzzing confusion*. New York: Oxford University Press.

[72] Weinstein, N . D . (1980): "Unrealistic Optimism About Future Life Events," *Journal of Personality and Social Psychology,* 39(5), 806—820.

[73] Warren, S. (1992): "Lower threshold for referral for psychiatric treatment for adopted adolescents." *Journal of the American Academy of Child and Adolescent Psychiatry* 31 n. 3.

[74] Wood, D., Bruner, J.S., and Ross, G. (1976): "The role of tutoring in problem-solving," *Journal of Child Psychology and Psychiatry*, **17**, 89-100.

[75] Wood, D., and Middleton, D. (1975): “A study of assisted problem-solving,” *British Journal of Psychology*, **66**, 181-191.

[76] Wood, D., Wood, H., and Middleton, D. (1978): “An experimental evaluation of four face-to-face teaching strategies,” *International Journal of Behavioral Development*, **2**, 131-147.

[77] Zentall, Thomas R. and Bennet G. Galef JR. (1988) *Social Learning: Psychological and Biological Perspectives.* Lawrence Erlbaum Associates, Publishers.

B Online Appendix

B.0.1 Proof of Proposition C.4 and Corollary C.5.

Begin by writing the parent's objective function in Eq. (34) in a way that is consistent with both the model of Sec. 2.1 and Sec. C.3. In both cases, the child's guess at time 2 takes the form $w(X_1+\bar{a})+(1-w)M_0$: if the child is unaware of the parent's action, then $w=\frac{p_X}{p_0+p_X}$, whereas, under the textbook equilibrium assumption, Lemma C.3 implies that $w=\frac{p_X(1-\gamma^e)}{p_0+p_X(1-\gamma^e)^2}$. We get

$$\mathrm{E}\left[(X_1+\bar{a}-M_0)^2+\delta[X_2-(1-w)M_0-w(X_1+\bar{a})]^2|M\right]. \tag{26}$$

It is also convenient to calculate the expectation: adding and subtracting M inside the two squared-loss terms, we get

$$\begin{aligned}
&\mathrm{E}\Big[(X_1+\bar{a}-M)+(M-M_0)]^2+ \qquad (27)\\
&+\delta[(X_2-M)+(1-w)(M-M_0)+w(M-X_1-\bar{a})]^2|M\Big]=\\
&=\bar{a}^2+\frac{1}{p_X}+(M_0-M)^2+2\bar{a}(M-M_0)+\\
&+\delta\left\{\frac{1}{p_X}+(1-w)^2(M_0-M)^2+w^2\left(\frac{1}{p_X}+\bar{a}^2\right)+2w(1-w)(M_0-M)\bar{a}\right\}=\\
&=(1+\delta w^2)\bar{a}^2-2[1-\delta w(1-w)](M_0-M)\bar{a}+ \qquad (28)\\
&+\delta\left\{(1-w)^2(M_0-M)^2+w^2\frac{1}{p_X}\right\}+\frac{1+\delta}{p_X}+(M_0-M)^2.
\end{aligned}$$

Differentiating with respect to $\bar{a}$ and rearranging terms yields

$$\bar{a}=\frac{1-\delta(1-w)w}{1+\delta w^2}(M_0-M), \tag{29}$$

and it is also immediate to verify that the second derivative of Eq. (28) is strictly positive.

For $w=\frac{p_X}{p_0+p_X}$, we obtain the optimal intensity $\gamma_T=\gamma_1$ in Eq. (11) of Thm. 2.1 (where $B_{T+1}=B_L=B_2=1$).

In the textbook-equilibrium case, Eq. (29) instead implies that a necessary and sufficient condition for a linear equilibrium with intensity γ^e is $\gamma^e=\frac{1-\delta(1-w)w}{1+\delta w^2}$. This can be rewritten as follows: notice that $[p_0+p_X(1-\gamma^e)^2]^2w(1-w)=p_X(1-\gamma^e)[p_0-p_X\gamma^e(1-\gamma^e)]$; then the required condition is

$$\gamma^e=\frac{[p_0+p_X(1-\gamma^e)^2]^2-\delta p_X(1-\gamma^e)[p_0-p_X\gamma^e(1-\gamma^e)]}{[p_0+p_X(1-\gamma^e)^2]^2+\delta p_X^2(1-\gamma^e)^2}.$$

Multiplying by the denominator, which is strictly positive, and rearranging yields

$$[p_0 + p_X(1-\gamma^e)^2]^2(1-\gamma^e) - \delta p_X(1-\gamma^e)[p_0 - p_X\gamma^e(1-\gamma^e) + \gamma^e p_X(1-\gamma^e)] = 0$$

which can be rewritten as

$$(1-\gamma^e)\Big\{[p_0 + p_X(1-\gamma^e)^2]^2 - \delta p_0 p_X\Big\} = 0.$$

Therefore, one solution is always $\gamma^e = 1$. To find other solutions, assume $\gamma^e \neq 1$, divide by $1-\gamma^e$ and rearrange:

$$p_X(1-\gamma^e)^2 = \sqrt{\delta p_0 p_X} - p_0.$$

Since $p_0 > 0$, the rhs is positive provided $\delta p_X > p_0$. If this condition holds, then the above holds if and only if

$$1-\gamma^e = \pm\sqrt{\frac{\sqrt{\delta p_0 p_X} - p_0}{p_X}}.$$

Note also that, if $\delta p_X > p_0$, then $\sqrt{\delta p_0 p_X} - p_0 < \delta p_X - p_0 < p_X$, which implies that the argument of the outer root is strictly between 0 and 1.

Choosing the positive sign in the above equation yields the expression for γ^e in Prop. C.4, henceforth denoted γ^+; by the argument just given, $\gamma^+ \in (0,1)$; Choosing the negative sign instead yields $\gamma^- \equiv 2 - \gamma^+$.

To complete the proof of Cor. C.5, it is convenient to express the minimized value of the parent's objective function by substituting for $\bar{a}$ in Eq. (28) using Eq. (29); we get

$$\begin{aligned}
&\frac{[1-\delta(1-w)w]^2}{1+\delta w^2}(M_0-M)^2 - 2\frac{[1-\delta w(1-w)]^2}{1+\delta w^2}(M_0-M)^2 + \\
&+\delta\left\{(1-w)^2(M_0-M)^2 + w^2\frac{1}{p_X}\right\} + \frac{1+\delta}{p_X} + (M_0-M)^2 = \\
&= \frac{-[1-\delta w(1-w)]^2 + \delta(1-w)^2 + \delta^2(1-w)^2w^2 + 1 + \delta w^2}{1+\delta w^2}(M_0-M)^2 + \frac{1+\delta(1+w^2)}{p_X} = \\
&= \frac{2\delta w(1-w) + \delta(1-w)^2 + \delta w^2}{1+\delta w^2}(M_0-M)^2 + \frac{1+\delta(1+w^2)}{p_X} = \\
&= \frac{\delta(1-w)[2w+1-w] + \delta w^2}{1+\delta w^2}(M_0-M)^2 + \frac{1+\delta(1+w^2)}{p_X} = \\
&= \frac{\delta}{1+\delta w^2}(M_0-M)^2 + \frac{1+\delta(1+w^2)}{p_X};
\end{aligned} \tag{30}$$

since $w^2 = \frac{p_X^2(1-\gamma^e)^2}{[p_0+p_X(1-\gamma^e)^2]^2}$ and $(1-\gamma^-)^2 = (\gamma^+-1)^2 = (1-\gamma^+)^2$, it follows that the payoff to the parent in the two equilibria with intensity coefficients γ^+ and γ^- is the same.

Finally, suppose that $\bar{a} = \gamma^e(M_0 - M)$ is an equlibrium parental intervention, with $b_1 = M_0$ and b_2 as in Lemma C.3, and fix $k \in \mathbb{R}$. Then let $\bar{a}' = \bar{a} + k$, $b_1' = b_1 + k$ and $b_2' = b_2$; clearly, (a, b) constitutes an equilibrium in which the child "filters out" the additive constant k both in her time-1 guess and in her updating rule. It is immediate to verify that this new equilibrium is payoff-equivalent to the original one.

C OMITTED: Extensions

C.1 Costly Parental Intervention

So far we have assumed that the parent's payoff per period coincides with that of the child for any given realization of uncertainty X_t and any choice of actions $\bar{a}_t$ and $\bar{b}_t$. In particular, this implies that the only "cost" of parental intervention is via the induced distortion of the child's beliefs.

We now extend our basic model and introduce an explicit cost of parental intervention. For analytical tractability, we assume that this additional cost is quadratic in the parent's action $\bar{a}_t$. Specifically, the *parent's* loss at time t, given X_t, $\bar{a}_t$ and $\bar{b}_t$, is now

$$(X_t + \bar{a}_t - \bar{b}_t)^2 + C\bar{a}_t^2 \tag{31}$$

where $C \geq 0$. Our main result, Theorem 2.1, admits a straightforward generalization to this environment.

Theorem C.1 *Under costly parenting, the optimal parenting policy $a = (a_1, \ldots, a_T) \in \mathcal{A}$ is $a_t = \gamma_t(M_{t-1}^a - M)$, where*

$$\gamma_t = \frac{1 - \delta B_{t+1} \frac{p_X}{p_t} \frac{p_{t-1}}{p_t}}{1 + C + \delta B_{t+1} \left(\frac{p_X}{p_t}\right)^2}, \quad B_{T+1} = \sum_{\tau=1}^{L-T} \delta^{\tau-1} \left(\frac{p_T}{p_{T+\tau-1}}\right)^2, \quad B_t = \frac{\delta B_{t+1}}{1 + \delta B_{t+1} \left(\frac{p_X}{p_t}\right)^2}. \tag{32}$$

Furthermore, $\gamma_t \in (0, 1)$.

Notice that the optimal parenting policy has exactly the same features as in the costless case; the only difference is the denominator of the intensity of intervention γ_t, which now includes the cost parameter C. Thus, it is still the case that partial sheltering occurs at the optimum, so our interpretation of the evidence from behavioral genetics (Sec. 3 is preserved under costly parenting.

C.2 Simultaneous learning in the Gaussian framework

Another interesting extension is *simultaneous learning.* Children typically engage in multiple simultaneous learning processes; furthermore, experimental evidence points to *complementarities* between different aspects of cognitive development.[30] It is then natural to ask how a supervisor might take advantage of these complementarities to resolve the sheltering/learning trade-off, and more generally aid in the development process.

The Gaussian framework can be readily extended to handle learning of multiple tasks. As we shall presently demonstrate, the main technical finding is that the parent's intervention in any given task is, in general, a linear function of the child's bias in *all* tasks. Preliminary numerical exploration suggests than the parent may take advantage of this in interesting and somewhat unexpected ways; this invites further investigation.

We continue to assume that the child lives for $L > 1$ periods, and that teaching terminates at time $T \in \{1, \ldots, L-1\}$ (i.e. there is at least one teaching period).

The child must now learn how to perform n tasks. Extending the notation of Section 2.1, let $M \in \mathbb{R}^n$ denote the vector describing the correct way to perform each task on average and; for every $t = 1, \ldots, T$, let X_t denote the n-dimensional random vector that describes the correct way to perform each task at time t.

The parent knows the value of M; the child has a multinormal prior over its possible values, with precision matrix p_0 and mean vector M_0. The parent and child both believe that, conditional upon M, each X_t is a multinormal vector with mean M and precision matrix p_X. Thus,

$$X_t | M \sim N\left(M, p_X^{-1}\right) \quad \text{and} \quad M \sim N\left(M_0, p_0^{-1}\right). \tag{33}$$

The parent's and child's actions at time t, denoted by $\bar{a}_t$ and $\bar{b}_t$ respectively, are also n-dimensional real vectors. The child's loss at time t is given by

$$(X_t + \bar{a}_t - \bar{b}_t)'(X_t + \bar{a}_t - \bar{b}_t).$$

We can also incorporate a quadratic cost-of-parenting term, extending the model in Sec. C.1. Thus, the parent's loss at time t is

$$(X_t + \bar{a}_t - \bar{b}_t)'(X_t + \bar{a}_t - \bar{b}_t) + \bar{a}_t' C \bar{a}_t,$$

where C is an $n \times n$ symmetric matrix.

By analogy with the single-task case, a parenting policy is a tuple $(a_1, \ldots, a_T)$ such that $a_t : \mathbb{R}^n \times \mathbb{R}^{n \cdot (t-1)} \to \mathbb{R}^n$ for each t. That is, the parent's intervention at time t depends

[30] For instance, Gopnick and Meltzoff [40] report that children who use more names are more likely to exhaustively sort objects into different categories (see also Waxman [71]).

upon the value of the vector M and the realizations of the random vectors $X_1, \ldots, X_{t-1}$. Similarly, a policy for the child is a tuple $(b_1, \ldots, b_L)$ such that, for every t, $b_t : \mathbb{R}^{n\cdot(t-1)} \to \mathbb{R}^n$. We continue to denote the set of policies for the parent and the child by $\mathcal{A}$ and $\mathcal{B}$ respectively.

Finally, in accordance with the assumptions discussed in the main text, given a parenting policy $(a_1, \ldots, a_T) \in \mathcal{A}$, the child's beliefs about M at the end of each time period $t = 0, 1, \ldots, L$ are characterized by the precision matrix $p_t \equiv p_0 + tp_X$ and the mean vector

$$M_t^a = \begin{cases} p_t^{-1}\left[p_0M_0 + p_X \sum_{s=1}^{t}[X_s + a_s(M, X_1, \ldots, X_{s-1})]\right] & t \leq T \\ p_t^{-1}\left[p_0M_0 + p_X \sum_{s=1}^{T}[X_s + a_s(M, X_1, \ldots, X_{s-1})] + p_X \sum_{s=T+1}^{t} X_s\right] & t > T \end{cases}$$

Also, the optimal policy $b^a = (b_1^a, \ldots, b_L^a) \in \mathcal{B}$ for the child, given the parent's policy $a = (a_1, \ldots, a_T) \in \mathcal{A}$, satisfies $b_t^a = M_{t-1}^a$ for all $t = 1, \ldots, L$. Again, these are immediate extensions of the corresponding results for the single-task case.

The parent still minimizes her discounted expected loss. Theorem 2.1 can then be generalized as follows.

Theorem C.2 *For $t = 1, \ldots, T$, the optimal parenting policy $a = (a_1, \ldots, a_T) \in \mathcal{A}$ satisfies $a_t = \Gamma_t \cdot (M_{t-1}^a - M)$, where Γ_t is inductively defined as follows:*

$$\begin{aligned} B_{T+1} &= \sum_{\tau=1}^{L-T} \delta^{\tau-1} p_t p_{t+\tau-1}^{-1} p_{t+\tau-1}^{-1} p_t, \\ \Gamma_t &= [I + C + \delta p_X p_t^{-1} B_{t+1} p_t^{-1} p_X]^{-1} [I - \delta p_X p_t^{-1} B_{t+1} p_t^{-1} p_{t-1}], \\ B_t &= (I - \Gamma_t)'(I - \Gamma_t) + \delta(p_{t-1} + p_X\Gamma_t)' p_t^{-1} B_{t+1} p_t^{-1} (p_{t-1} + p_X \Gamma_t). \end{aligned}$$

C.3 "Textbook" Equilibrium analysis

With regards to equilibrium analysis of the "Hand-Holding" models, three main findings can be mentioned; we focus on the single-task environment for simplicity.

First, if parental intervention is intrinsically costless (as in our main model), then for any choice of the parameters, the following profile of strategies is a (perfect Bayesian) equilibrium: the parent chooses $a_t = M_0 - M$ for $t = 1, \ldots, T$; the child guesses M_0 in the first T periods, disregarding the realization of $X_t + a_t$, and then proceeds as in the unsupervised case from time $T + 1$ onwards. In this equilibrium the parent opts for a *fully sheltering* strategy, which however prevents learning in the first T periods: the child is aware that he is essentially observing draws from a normal distribution with mean $M + a_t = M_0$ in each period, regardless of the realized value of M, and hence he simply discards these observations. Thus, in this equilibrium, the child does not learn anything while being supervised by the parent. However, it is easy to see that this particular equilibrium disappears as soon as costs are introduced.

Second, under natural assumptions on parameters, there exists a partially sheltering equilibrium in which the parent's strategy is linear in the child's bias, as in Theorem 2.1. In particular, it is characterized by a system of equations similar to Eq. (11), except that the quantities p_t and B_t are replaced by functions of all the equilibrium values of $\gamma_1 \dots \gamma_T$.

Third, simple two-period examples demonstrate that the child's welfare can be *lower or higher* in equilibrium than under the assumptions about beliefs stated above. Intuitively, the child can make better use of the information he receives (because he can partially filter out the parent's intervention); on the other hand, this induces the parent to shelter more than she would under the original assumptions about beliefs.

To illustrate these points, we solve a two-period, equilibrium version of the basic model in Sec. 2.1 (this can be extended to arbitrarily many periods). Since a policy for the parent consists of a single function of the true mean M, we drop time indices and denote it simply by $a : \mathbb{R} \to \mathbb{R}$. Equilibrium quantities will be denoted by the subscript "e".

Begin by conjecturing a *linear equilibrium*: that is, a Bayesian Nash equilibrium where the parent's equilibrium policy is linear in the child's initial bias: $a^e(M) = \gamma^e(M_0 - M)$, where $\gamma^e \in \mathbb{R}$ is the equilibrium intensity of intervention (again dropping the time index). Then, in equilibrium, the child knows that her time-1 observation is a realization of $X_1 + a(M) = X_1 + \gamma^e(M_0 - M)$, and not just X_1; thus, she updates accordingly. The following Lemma provides the details.

Lemma C.3 *For all* $\gamma \in \mathbb{R}$,

$$\mathrm{E}\left[M | X_1 + \gamma(M_0 - M) = x\right] = \frac{p_0 - p_X\gamma(1-\gamma)}{p_0 + p_X(1-\gamma)^2} M_0 + \frac{p_X(1-\gamma)}{p_0 + p_X(1-\gamma)^2} x.$$

In a linear equilibrium with intensity of intervention γ^e, the child's policy $b^e = (b^e_1, b^e_2)$ must then satisfy

$$b^e_1 = M_0,^{31} \quad b^e_2(x) = \mathrm{E}[M | X_1 + \gamma^e(M_0 - M) = x] \; \forall x \in \mathbb{R}.$$

Consequently, a necessary and sufficient condition for the existence of such an equilibrium

[31] In any equilibrium, the child's guess at time 1 only influences her loss at time 1, because the parent must choose an action $a(M)$ without observing b^e_1. Thus, in particular, in a linear equilibrium the child's optimal choice is $b^e_1 = \mathrm{E}\left[X_1 + \gamma^e(M_0 - M)\right] = M_0$.

is that, for every $M \in \mathbb{R}$,

$$\begin{aligned}\gamma^e(M_0 - M) \quad &\in \quad \arg\max_{\bar{a}\in\mathbb{R}} \mathrm{E}\Bigg[(X_1 + \bar{a} - M_0)^2 + \\ &+ \quad \delta\left(X_2 - \frac{p_0 - p_X\gamma^e(1-\gamma^e)}{p_0 + p_X(1-\gamma^e)^2}M_0 - \frac{p_X(1-\gamma^e)}{p_0 + p_X(1-\gamma^e)^2}(X_1 + \bar{a})\right)^2 \Bigg| M\Bigg]\end{aligned} \tag{34}$$

The following Proposition and Corollary characterize the linear equilibria of this game.

Proposition C.4 *For all parameter values, there is a linear equilibrium with* $\gamma^e = 1$. *Furthermore, if* $\delta p_X > p_0$, *there is a linear equilibrium with*

$$\gamma^e = \hat{\gamma} \equiv 1 - \sqrt{\frac{\sqrt{\delta p_0 p_X} - p_0}{p_X}} \in (0,1).$$

Corollary C.5 *If there is a linear equilibrium with* $\gamma^e = \hat{\gamma}$, *then there is also a payoff-equivalent linear equilibrium with* $\gamma^e = 2 - \hat{\gamma}$; *there are no other linear equilibria. Furthermore, if* $\gamma^e \in \{\hat{\gamma}, 2-\hat{\gamma}, 1\}$ *is the intensity of intervention in a linear equilibrium, then for every* $k \in \mathbb{R}$ *there is a payoff-equivalent equilibrium with* $a(M) = k + \gamma^e(M_0 - M)$.

Thus, as noted above, under the condition $\delta p_X > p_0$, one equilibrium exhibits the essential features of the solution we consider in Sec. 2: the parent's policy is linear in the child's bias, and in particular the intensity coefficient lies between zero and one. This coefficient is also decreasing in δ, as in Theorem 2.1. Such an equilibrium exists, provided the child is not too slow to learn relative to the complexity of the task: otherwise, the parent fully shelters ($\gamma^e = 1$) in any linear equilibrium.

Also, adding a constant k to the parent's action a leads to a formally distinct but payoff-equivalent equilibrium: the child merely adjusts her guesses b_0 and b_1 so as to offset the constant k. Finally, it turns out that, if there is an equilibrium with $\gamma^e \in (0,1)$ as in Prop. C.4, then there is another linear equilibrium with intensity coefficient equal to $2 - \gamma^e$: intuitively, this is the "mirror image" of the original linear equilibrium; the child can also adjust her learning rule accordingly, and again one obtains a payoff-equivalent equilibrium. The above Corollary summarizes these facts.

Table 2 below shows that, depending on parameter values, the child's welfare, as perceived by the informed parent, may be either higher or lower under textbook equilibrium behavior. The figures in Table 2 suggest that, if the child's initial bias is large, then for relatively low values of p_X the child's expected loss is higher in a linear equilibrium than under the assumption that the child is unaware of the parent's intervention, whereas the opposite is true for relatively high values of p_X. However, this ranking is reversed if the child's bias is small. Additional numerical experimentation seems to confirm this pattern. Finally, the expected loss in a linear equilibrium with $\gamma^e \in (0,1)$ may be either higher or lower than in the full-sheltering equilibrium with $\gamma^e = 1$; as one might expect, the full-sheltering equilibrium is superior for if the child's initial bias is small.

δ	p_0	p_X	$M_0 - M$	γ	Loss	γ^e	Eqm Loss	Eqm loss ($\gamma^e = 1$)
0.9	1	2	2	0.5714	3.7214	0.5867	3.8041	4.55
0.9	1	5	2	0.5385	2.7204	0.5264	2.3013	3.98
0.9	1	2	0.5	0.5714	1.3107	0.5867	1.2885	1.175
0.9	1	5	0.5	0.5385	0.6435	0.5264	0.7103	0.605

Table 2: Textbook Equilibrium and Child Welfare

D OMITTED: Discussion

D.1 Additional Related Literature

Influencing other agents' beliefs is a recurring theme in information economics and game theory. The main contrast between our model and most contributions on communication in models of asymmetric information is the following. In these models, the presence of conflicting incentives either makes it impossible to credibly communicate all the information or introduces the necessity of distorting actions in order to make communication credible. In our model on the other hand, although incentives are aligned, information cannot be directly communicated but must be transmitted through a (possibly costly) manipulation of the learning process. The informed agent in our models faces very different incentives from those in standard models of asymmetric information.

A "teaching" metaphor is sometimes employed to describe strategies in the literature on learning in games: cf. e.g. Kalai and Lehrer [48], Section 2.4. However, this seems mainly an expository device to clarify features of certain strategies; players do not have any private information to convey (besides their future dispositions to play). In the absence of the "teacher," the "learner" would have nothing to learn, which is definitely not the case in supervised learning models.

The literature on social learning is also related. In particular, Smith and Sorensen [65] analyze optimal experimentation and its interaction with informational herding. Although teaching is not the focus of their paper, these authors provide a two-period example featuring a "professor" who takes the more informative of two possible actions so as to maximize the "student's" expected payoff. This intriguing example may point in the direction of a full-blown model of *imitative* learning, or "teaching by doing" on the parent's part; by comparison, our concern in this paper is "learning by doing" on the child's part, under the supervision of the parent.

Benabou and Tirole [10] consider the interplay between the provision of incentives and the self-confidence of an agent: provision of explicit incentives by an informed principal (e.g. a parent) may be a negative signal of the ability of an agent (the child). They show that in some circumstances, explicit incentives can be counterproductive. Other authors (e.g., Rubinstein [61]) and Brunnemeier and Parker [22]) have investigated how beliefs may be optimally chosen by a 'principal' to maximize the payoffs of an 'agent'. In these settings, optimistic beliefs may be advantageous.

Ettinger and Jehiel [35] propose a model of deception in which "manipulative" agents influence the beliefs of less sophisticated, "deceived" agents. In particular, less sophisticated agents make inferences on the basis of coarser information about their opponent than is available in the game—an assumption consistent with the so-called "fundamental attribution error." This bounded-rationality assumption leads to the possibility of deception. It should be noted, however, that no deception is possible if both agents are fully rational; by way of contrast, as we show in Sec. C.3, the main qualitative features of our analysis are consistent with textbook-equilibrium analysis.

Finally, the transmission and adoption of cultural traits had been developed by, among others, Cavalli Sforza-Feldman [25], and Boyd-Richerson [20]. In most of these literature, the transmission is independent of any of the choices of the parents. Bisin and Verdier [13], [14], and Bisin, Topa, and Verdier [15] develop models in which cultural transmission occurs as the result of the socialization effort endogenously chosen by the parent. However, the effectiveness of socialization is modeled as a reduced-form function of some effort measure.

D.2 Applications

The approach outlined in this paper may be useful to address broader issues pertaining to learning in a multi-person environment. First, at an abstract level, **schooling** may be viewed as supervising more than one simultaneous learning process, with only limited opportunities to target interventions to individual learners. For instance, in the setting of §2.1, assume that there n 'students' with different priors about M, and one 'instructor'. At every time t, each student must guess the realization of the r.v. X_t; the instructor must choose the same level of intervention a_t for all students. If $b_{t,i}$ is student i's guess at time t, his payoff is $-(X_t + a_t - b_{t,i})^2$. If students do not observe each other's actions, the model is an n-fold repetition of the model in Section 2.1, except that the instructor's choice of a_t must be the same in every "copy" of the problem. We can then analyze the impact of different objective functions for the instructor on (i) the distribution of ability in the population of students, and (ii) the ability of a single student, given the distribution of other students' abilities.

A more interesting model of **peer effects** may plausibly be obtained by allowing students to observe each other's guesses. This would likely be beneficial for "low-ability" students; however, it might be the case that high-ability students are better off than they would be if actions are unobservable. Intuitively, in a very heterogeneous class, the instructor's choice of a_t is likely to be significantly different from the choice she would make if she was interacting with a single high-ability student; but, as low-ability students improve due to the presence of high-ability peers, the class becomes less heterogeneous, and as a consequence even high-ability students may benefit.