

Parental Guidance and Supervised Learning

Alessandro Lizzeri

Marciano Siniscalchi*

June 27, 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*.

*We thank Daron Acemoglu, Joe Altonji, Greg Duncan and Jim Heckman, as well as Ed Glaeser and three anonymous referees. Financial support from NSF Grants SES-0452317 (Lizzeri) and SES-0453088 (Siniscalchi) is gratefully acknowledged. Aaron Hauptmann provided excellent research assistance.

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 away from nurture, downplaying parental influence.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, once one controls for genetic factors.³ At the same time, parents spend great energy and resources in attempts to affect long-run outcomes for their children. How can we reconcile evidence suggesting that parents do not matter with common perceptions and practices that suggest that, in fact, they do?

This paper provides a model of optimal parenting that can help resolve this contradiction. In the model we propose, parents have significant effects on the characteristics of their children. Yet, the model predicts a distribution of outcomes in the population of children that is consistent with the data from behavioral genetics. This calls into question the interpretation of such data that has been made in that literature. Furthermore, our model allows an interpretation of the features of the environment, of the characteristics of the parents, and of the interactions between parents and children that are responsible for these patterns in the data. In contrast, the standard analytical framework of behavioral genetics is essentially a statistical model that does not naturally lend itself to such interpretations or to policy (thought) experiments. Finally, our model provides a novel interpretation of measures of heritability that have received considerable attention in the social sciences.⁴

To provide some background for our analysis, the behavioral–genetics literature starts with a criticism of much of traditional developmental psychology for failing to recognize that ‘nature’ could be behind correlations between parenting styles and children’s outcomes because of shared genes between parents and their biological children.⁵ Behavioral geneticists then attempt to isolate the effects of the genes through several complementary approaches. The more direct one is to compare twins raised together by their biological parents with twins raised apart by different adoptive families.⁶ A consistent finding of several studies is that *twins reared together are just as similar as twins reared apart*. In fact, some studies even find that twins reared

¹For some history on this debate, see Pinker [31].

²For an overview of this literature, see Plomin et al. [33] or Reiss et al [34]. Some of the findings of this literature have been popularized by Harris [27], [28], Pinker [31] and Ridley [35].

³See for instance Harris [27], [28], Pinker [31] and Ridley [35].

⁴See for instance the debate surrounding the controversial book by Herrnstein and Murray [29].

⁵This branch of developmental psychology is known as socialization research. For a survey, see for instance Collins et al. [18] and Demo and Cox [20].

⁶There 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. [14].

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.⁷

Our model of *supervised learning* permits a different interpretation of the same data. In our framework, children’s characteristics evolve through interactions with both the environment and their parents; every parent faces a basic trade-off between *sheltering* the child from the consequences of his mistakes, and allowing him to *learn from experience*.

The key ingredients of our model are as follows. 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, and *learns by doing*: at the end of each period, he receives a signal about the quality of his performance. However, *learning is costly*: his payoff is lower the worse his performance. The parent has better information than the child about the correct way to perform the task, and can take actions that simultaneously modify (typically increase) the child’s per-period payoff and distort (typically bias) the child’s signal about his performance.

We characterize the optimal parenting policy and its comparative-statics properties. In particular, 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 smaller cost than if he were on his own. We also investigate the dynamics and comparative statics of parental intervention: see Sec. 2.1 for details.

These key implications of our model allow us to tackle the nature-nurture debate. We can think of the child’s initial abilities as being genetically determined; similarly, the parent’s abilities are determined by her initial genetic endowment, as well as a prior learning phase. The parent partially shelters the child from the consequences of his mistakes, as the parent herself perceives them in light of her own ability; this immediately implies that, in our model, children of “better” parents have better outcomes, even controlling for the child’s genetic endowment. On the other hand, parental intervention is *genetically mediated*: it responds to the child’s initial abilities. We show that this has an important implication: adoptive parents provide more sheltering *on average* than biological parents.⁸ Hence, on average, adoptive children are less exposed to environmental influences; thus, genetic similarity in initial abilities induces greater correlation in the outcomes for adoptive children. This key implication enables our model to

⁷This broad conclusion is subject to two qualifications: first, there is evidence that improvements in the family environment have positive effects on children’s cognitive ability, if one restricts attention to families of low socioeconomic status (Turkheimer et al. [42]). 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 [34]). Thus, family environment may have significant effects, even though its impact is “genetically mediated.”

⁸See sec. 3 for details, and a discussion of the evidence supporting this finding.

match the empirical findings from behavioral genetics for a broad and easily interpretable range of parameter choices. We also show that, in the optimal policy, 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 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 [5]); a particularly relevant contribution is the recent paper by Cunha and Heckman [19], which analyzes the optimal life-cycle profile of investments in children in a model with dynamic complementarities. This approach enables the authors to capture some important facts about child development. Akabayashi [1] provides a model of a child's human capital formation that can explain child maltreatment. Weinberg [46] obtains a positive relation between parents' income and children's outcomes via a model where parents use pecuniary incentives and corporal punishment to affect the behavior of children.

The literature on the consequences of parenting is virtually unanimous in recognizing that parental support is essential for functional development in extreme situations.¹⁰ The nature-nurture debate pertains mainly to differences within the 'normal' range of variation in parenting. A sizable literature across the social sciences argues that these differences can have important effects. An enormous literature in developmental psychology addresses the effects of parental care on the development of children.¹¹ 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 attain-

⁹In 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 [23]. Camerer [15], chapter 6 discusses both theory and experiments.

¹⁰Harlow and coauthors (cf. [25], [26]) 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 (Chisholm [17]).

¹¹It is impossible to be exhaustive in providing references. See Shonkoff and Deborah Phillips eds. [39] for a recent overview of the field with respect to early childhood development.

ment 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.¹² Recent contributions by Sacerdote [37, 38], Bjorklund, Lindahl and Plug [11] and others are methodologically closer to the behavioral-genetics literature. For instance, Sacerdote [38] 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. [11], and Bjorklund et al. [12] 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.

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 $t = 1, \dots, L$ is represented by i.i.d. normal random variables X_1, \dots, 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*. Alternatively, one can think of the sequence of events

¹²See e.g. Carneiro, Heckman, and Masterov [16] and Todd and Wolpin [41]. 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, [4]; Desai et. al. [21]; Belsky and Eggebeen [6], Bernal [8]), others that it is beneficial (Vandell and Ramanan [43]). See also the debate on the effects of family size and birth order (e.g., Black, Devereux, and Salvanes) [13].

as follows: first the problem arises, then the child makes a decision to confront the problem, and then, after observing the child’s choice, the parent chooses a corrective action.

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.

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, \dots, 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.¹³

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 our setting this takes a convenient form:

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

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 $x_1 + \bar{a}_1, \dots, x_{t-1} +$

¹³This 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.

\bar{a}_{t-1} , the child’s optimal action is his conditional expectation of X_t :

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

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 $M_0^a = M_0$.

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:

$$E[(X_t + a_t - M_{t-1}^a)^2 | X_1, \dots, X_{t-1}, M] = p_X^{-1} + (M + a_t - M_{t-1}^a)^2. \quad (3)$$

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: its solution involves a simple adaptive rule. Our main findings can be generalized to a suitable class of non-Bayesian-rational adaptive learning rules.

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 3.¹⁴

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: $a_t = M_{t-1}^a - M$ (a.k.a. “The Italian Mom”). This policy minimizes the per-period loss at times $t = 1, \dots, T$: this can be seen from Eq. (3). Intuitively, recall that the child’s

¹⁴Indeed, it is well known at least since the work of Bandura and coauthors [3, 2] 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 [44], emphasizes the “social” aspects of child development.

choice at time t is $b_t = M_{t-1}^a$, and the loss is $(X_t + a_t - b_t)$: thus, by choosing $a_t = M_{t-1}^a - M$, 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_{t-1}^a 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: $a_t = \frac{p_0 + (t-1)p_X}{p_X}(M - M_{t-1}^a)$. This policy ensures that, at the end of time t (i.e. after observing X_t), the child’s posterior M_t^a 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. Notice that, as time goes by, intervention becomes more and more severe; this is because, under Bayesian updating, the child’s posterior precision also increases (at time t , it equals $p_0 + t p_X$), so it becomes harder to “convince” him that he is wrong.

Thus, the present framework allows for a range of qualitatively very different parenting strategies involving positive as well as negative learning effects. 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: $a_t = \gamma_t(M_{t-1}^a - M)$. 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 (and robust) conclusion of Theorem 2.1 is the finding that optimal parenting entails *partial sheltering*: the intensity of intervention γ_t lies in $(0, 1)$.¹⁵ 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 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”). In particular,

¹⁵We have verified numerically that, in a two-period model, partial sheltering is optimal for a wide range of loss functions of the form $|X_t + \bar{a}_t - \bar{b}_t|^r$, with r ranging from 0^+ to substantially above 2.

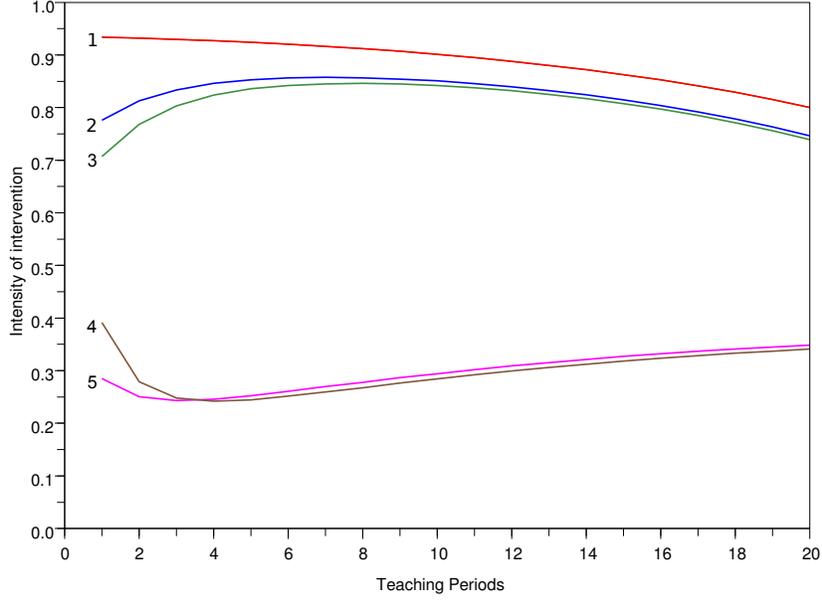


Figure 1: Intensity of intervention for different values of p_0 and p_X .

$\gamma_t = \mu_t \gamma_t^{\text{FS}} + (1 - \mu_t) \gamma_t^{\text{BC}}$, 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.

Label	Color	δ	p_0	p_X	Label	Color	δ	p_0	p_X
1	Red	0.9	1	0.1	4	Brown	0.99	0.1	1
2	Blue	0.9	1	1	5	Magenta	0.99	1	1
3	Green	0.9	0.1	1					

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

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**¹⁶ can be understood in terms of the decomposition of γ_t into a weighted average of the “Boot Camp” and “Full Sheltering” 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.¹⁷ 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.

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.

¹⁶Formal statements and proofs of the results referred to here can found in the Web Appendix available from the authors’ Web pages.

¹⁷Intuitively, 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.

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 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_{Z_0} . 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 fitness measure: agents who are closer to the true mean, namely, those who are more correct on average 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 $N(Z_0, p_{Z_0})$ 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 elaborate on this point in Sec. 3.4, and provide explicit calculations in Appendix A.2. Additionally, we 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 the prior means of two individuals’ play a crucial role in our analysis. *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 monozygotic twins are perfectly correlated; and so on.¹⁸

¹⁸We note that there may be some ambiguity as to the interpretation of the term ‘parent.’ We typically mean some

An convenient 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 Section 2. 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, \dots, X_t .

3.2 Behavioral Genetics: the ACE model

As was noted in the Introduction, the literature on behavioral genetics (BG henceforth) emphasizes the central role of an individual’s genes in determining a variety of traits such as IQ, introversion, neuroticism, social attitudes and many others. Behavioral geneticists build much of their analysis on the so-called “ACE model”; we now discuss a widely-used variant (cf. Plomin, et al. [33], 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 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*. 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. In the words of Goldberger [24], the factors A , C and E are 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

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, as well as the intensity of intervention γ_1 . Specifically,

$$\underbrace{M_1}_P = \underbrace{\frac{p_0}{p_0 + p_X} M_0}_A + \underbrace{\gamma_1 \frac{p_X}{p_0 + p_X} (M_0 - Z_0)}_C + \underbrace{\frac{p_X}{p_0 + p_X} X_1}_E, \quad (5)$$

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.¹⁹ 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.

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 [14], 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 (cf. Goldberger [24]).

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 envi-

¹⁹An 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.

ronment; 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. 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 because 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. In other words, our model can match this data, but is consistent with a more positive view of parenting. Towards the end of §3.4, we also argue that other natural alternative models, including ones in which parents directly communicate with children, cannot rationalize these findings. Our model also allows us to obtain additional predictions that have not yet been investigated in the BG literature.

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).

Appendix A.2 explicitly calculates $r^{MZT} - r^{MZA}$ and $r^{MZT} - r^{DZT}$; here, we focus on

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

We emphasize a key 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*. Thus, common rearing has a direct (and obvious) positive effect on phenotypic correlation, reflected in the first term in the r.h.s. of Eq. (6); however, our analysis uncovers a *compensating effect*. 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, for a broad range of sensible parameterizations. Thus,

it is useful to elaborate upon the intuition behind the negative effect of r_0 on the difference in phenotypic correlation. We shall then discuss our quantitative findings.

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 [30] 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 [45] 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.”

The fact that sheltering is greater for adoptive children implies 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;²⁰ this is captured in Eq. (6) by the negative coefficient of r_0 .

Other models can potentially 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.²¹ 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. This suggests that the effect outlined in this paper should be robust to *some* alternative modeling approaches.

However, it is useful to point out that other natural, alternative modeling approaches would *not* generate such an effect. For instance, consider a variant of the model developed in this paper in which parents simply communicate a noisy signal of their prior mean Z_0 to children, but do not intervene in their learning process. Such a model would differ from the reference

²⁰The 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_{Z_0} , the variance of the parent’s prior; see also Eq. (22) in Appendix A.2.

²¹Bisin and Verdier [9, 10] make such an assumption in a model of cultural transmission.

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 $r^{MZT} - r^{MZA}$ is small, unless one is willing to assume that communication is particularly ineffective.²² But, in a model of this kind, this would be tantamount to *assuming* directly that parents have limited influence on their children, as suggested by Harris [27] and other “extreme anti-nurturists.”

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,²³ 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_{Z_0} and correlation r_0 with M_0 (for twins reared by their biological parents) as functions of the remaining parameters. Appendix A.2 provides the details.

We set $\delta = (0.95)^{20}$, intuitively suggesting a 95% yearly discount factor and a teaching period lasting for 20 years.²⁴

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, it is easily verified 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.

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.

²²Also, regardless of the level of communication, such a model could not account for the finding that $r^{MZT} < r^{MZA}$ for some traits.

²³For 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.

²⁴Higher values of δ lead to even greater prevalence of the above correlational patterns. Also, the present calculations assume $L = 2$; however, a nearly indistinguishable diagram is obtained for higher values of L —for instance, consistently with Fig. 1, one could take $T = 1$ and $L = 5$, where each “period” corresponds to 20 years.

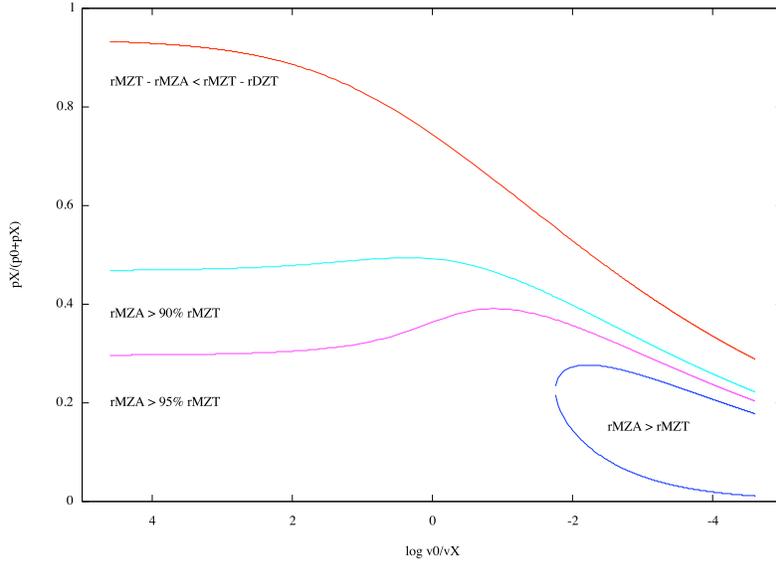


Figure 2: Twins reared together vs. Twins reared apart

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—hence, a relatively limited contribution of genetic factors to the phenotype. As Fig. 2 shows, whether, for instance, $r^{MZA} - r^{MZA}$ is smaller or larger than $r^{MZA} - r^{DZA}$ 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 $r^{MZA} - r^{MZA} < r^{MZA} - r^{DZA}$. Behavioral geneticists would interpret this as indicating that genetic factors are determinant; however, our model suggests that, in these circumstances, 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 [35, 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 Auxiliary predictions of the model

The observation that adoptive parents provide more sheltering have two further empirical implications in our model that have not been discussed in the BG literature.

Differences in phenotypic variance 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. The Appendix provides an explicit calculation.²⁵

Thus, our model provides a prediction on the relative dispersion of adopted vs. non-adopted children. This auxiliary prediction of our model *cannot* be obtained in the ACE framework: in the latter, phenotypic variances are the same in all sub-populations. We have surveyed several studies to verify whether this prediction is consistent with the data. The answer is largely positive, albeit not in every sample or for every outcome measure. For instance, in Bjorklund et al [12] (see Table 1), the variance of earnings and income is larger for adoptive children, while the variance of years of education is slightly smaller for adoptive children. We view this as good news for our stark model of parenting, since there are potential countervailing factors pushing the variances in the opposite direction. In particular, because the process of adoption is somewhat selective, the most extreme dysfunctional families are likely to be excluded, thereby reducing the population variance of adoptive parents; as the calculations in Appendix A.2 show, this in turn reduces the variance for adoptive children.

Correlation of MZ twins adopted together In the previous subsection we have discussed the contrast between the correlations of MZ twins reared together by their biological parents and those reared apart by different adoptive parents, because this is a focus of the BG literature. However, it is easy to show that our model predicts that the highest correlation among twins arises for MZ twins reared together in an adoptive family. The reason is that these individuals share the same parents, and at the same time they are sheltered more than children reared by biological families. When comparing MZA and MZT twins, these two forces push in opposite direction; however, they both lead to an increased correlation for twins reared in the same adoptive family.

3.6 Heritability Estimates

In the context of the ACE model, a key quantity of interest is **heritability**, defined as

$$\frac{\text{Var}[A]}{\text{Var}[P]} \equiv h^2. \tag{7}$$

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

²⁵It 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.

[14] reports values of heritability ranging from 69% to 78%. According to the above interpretation, these figures 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 [24]); nevertheless, heritability remains a central quantity of interest in BG, and one that is often invoked in various instances of the nature-nurture debate (e.g. Herrnstein and Murray [29]).

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 common 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. The reader is referred to Goldberger [24] and Plomin et al. [33] for rigorous calculations;²⁶ we emphasize that crucial independence assumptions are required in these derivations.

On the other hand, we show in Appendix A.2 that, in our model,

$$r^{MZA} = 2(r^{MZT} - r^{DZT}) = [(1 - p) + \gamma_1 p]^2 h^2.$$

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.

A Appendix

A.1 Theorem 2.1

²⁶For convenience, the Web Appendix available from the authors' Web pages also provides calculations.

Let $p_t = p_0 + t p_X$ for $t = 0, \dots, L$, and define the random variable

$$M_t^a = \begin{cases} \frac{p_0 M_0 + p_X \sum_{s=1}^t [X_s + a_s(M, X_1, \dots, X_{t-1})]}{p_t} & t \leq T \\ \frac{p_0 M_0 + p_X \sum_{s=1}^T [X_s + a_s(M, X_1, \dots, X_{t-1})] + p_X \sum_{s=T+1}^t X_s}{p_t} & t > T \end{cases} \quad (8)$$

The formal statement of Theorem 2.1 is as follows:

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

$$a_t = \gamma_t (M_{t-1}^a - M), \quad (9)$$

$$\text{where } \gamma_t = \frac{1 - \delta B_{t+1} \frac{p_X}{p_t} \frac{p_{t-1}}{p_t}}{1 + \delta B_{t+1} \left(\frac{p_X}{p_t}\right)^2}, \quad B_{T+1} = \sum_{\tau=1}^{L-T} \delta^{\tau-1} \left(\frac{p_T}{p_{T+\tau-1}}\right)^2, \quad B_t = \frac{\delta B_{t+1}}{1 + \delta B_{t+1} \left(\frac{p_X}{p_t}\right)^2}. \quad (10)$$

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

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

Proof. The fact that γ_t is decreasing in δ and L follows from a simple induction argument, by inspecting Eq. (10); we thus focus on the remaining statements. Let $a = (a_1, \dots, a_T) \in \mathcal{A}$ denote the parent's optimal policy; also, to simplify the notation, we write M_{t-1}^a 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. From Eq. (8), for all $t \geq T$ and $\tau \geq 0$,

$$M_{t+\tau} = \frac{1}{p_{t+\tau}} \left(p_t M_t + p_X \sum_{s=t+1}^{t+\tau} X_s \right), \quad (11)$$

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

$$\begin{aligned} X_{t+\tau} - M_{t+\tau-1} &= X_{t+\tau} - \frac{p_t}{p_{t+\tau-1}} M_t - \frac{p_X}{p_{t+\tau-1}} \sum_{s=t+1}^{t+\tau-1} X_s = \\ &= (X_{t+\tau} - M) - \frac{p_t}{p_{t+\tau-1}} (M_t - M) - \frac{p_X}{p_{t+\tau-1}} \sum_{s=t+1}^{t+\tau-1} (X_s - M); \end{aligned} \quad (12)$$

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 $\tau \geq 1$, the expected loss at time $t + \tau$ given the observed value of M_t and the true value M is

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

because X_1, \dots, X_L are i.i.d. $N(M, p_X^{-1})$ given M , so 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 $s \in \{t+1, \dots, t+\tau-1\}$, and $(X_s - M)(X_\sigma - M)$ for s, σ distinct in $\{t+1, \dots, t+\tau-1\}$, all have zero conditional expectation.

Hence, for every $t+1 \in \{T+1, \dots, L\}$, 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

$$A_{t+1} = \sum_{\tau=1}^{L-t} \delta^{\tau-1} \left[\frac{1}{p_X} + \frac{p_X}{p_{t+\tau-1}^2} (\tau-1) \right], \quad B_{t+1} = \sum_{\tau=1}^{L-t} \delta^{\tau-1} \left(\frac{p_t}{p_{t+\tau-1}} \right)^2. \quad (13)$$

Turn now to teaching periods $t \in \{1, \dots, T\}$. 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 $t = T, \dots, 1$, if $V_{t+1}(M_t, M) = A_{t+1} + B_{t+1}(M_t - M)^2$, 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. By Eq. (8), for every $t \leq T$,

$$M_t = \frac{p_0}{p_t} M_0 + \frac{p_X}{p_t} \sum_{s=1}^t (X_s + a_s) = \frac{p_{t-1}}{p_t} M_{t-1} + \frac{p_X}{p_t} (X_t + a_t). \quad (14)$$

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

$$\begin{aligned} V_t(M_{t-1}, M) &= \min_{\bar{a}} \mathbb{E} [(X_t + \bar{a} - M_{t-1})^2 | \mathcal{I}_t] + \\ &+ \delta \mathbb{E} \left[A_{t+1} + B_{t+1} \left(\frac{p_{t-1}}{p_t} M_{t-1} + \frac{p_X}{p_t} (X_t + \bar{a}) - M \right)^2 | \mathcal{I}_t \right]. \end{aligned} \quad (15)$$

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

$$\begin{aligned} 0 &= \bar{a} - (M_{t-1} - M) + \delta B_{t+1} \left(\frac{p_{t-1}}{p_t} M_{t-1} + \frac{p_X}{p_t} (M + a) - M \right) \frac{p_X}{p_t} \quad \text{and therefore} \\ a_t &= \frac{1 - \delta B_{t+1} \frac{p_X p_{t-1}}{p_t^2}}{1 + \delta B_{t+1} \left(\frac{p_X}{p_t} \right)^2} \cdot (M_{t-1} - M) \equiv \gamma_t (M_{t-1} - M). \end{aligned} \quad (16)$$

We now show that V_t can be expressed as a quadratic form in $(M_{t-1} - M)$. First, note 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. (15),

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

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

$$B_t = (1 - \gamma_t)^2 + \delta B_{t+1} \left(\frac{p_{t-1} + p_X \gamma_t}{p_t} \right)^2. \quad (17)$$

To show that B_t can be written as in Eq. (10), note that

$$1 - \gamma_t = \frac{1 + \delta B_{t+1} \left(\frac{p_X}{p_t} \right)^2 - 1 + \delta B_{t+1} \frac{p_X p_{t-1}}{p_t^2}}{1 + \delta B_{t+1} \left(\frac{p_X}{p_t} \right)^2} = \frac{\delta B_{t+1}}{1 + \delta B_{t+1} \left(\frac{p_X}{p_t} \right)^2} \frac{p_X}{p_t} \quad \text{and} \quad (18)$$

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

The expression for B_t in Eq. (10) now follows from Eqs. (17), (18), and (19).

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. (13), 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, by induction, let $t \in \{1, \dots, T\}$; assume that the claim is true for $t+1$. By Eqs. (18) and (10), $B_{t+1} = (1 - \gamma_{t+1}) \frac{p_{t+1}}{p_X}$, 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 $p_t = p_0 + t p_X \geq p_0 + p_X$ 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

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

and Eq. (10) 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 Supervised Learning and Twin Correlations

Recall that $p = \frac{p_X}{p_0 + p_X}$ 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

$$v_1 = [(1-p) + p\gamma_1]^2 v_0 + p^2 v_X + [p\gamma_1]^2 v_{Z_0} - 2p(1-p)r_0 \sqrt{v_0 v_{Z_0}}. \quad (20)$$

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

$$v_{1a} = [(1-p) + p\gamma_1]^2 v_0 + p^2 v_X + [p\gamma_1]^2 v_{Z_0}. \quad (21)$$

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 $M'_1 = [(1-p) + p\gamma_1]M_0 - p\gamma_1 Z_0 + pX'_1$, because these siblings have the same genetic endowment and the same parents. Thus

$$r^{MZT} = [(1-p) + p\gamma_1]^2 \frac{v_0}{v_1} + [p\gamma_1]^2 \frac{v_{Z0}}{v_1} - 2[(1-p) + p\gamma_1]p\gamma_1 r_0 \frac{\sqrt{v_0 v_{Z0}}}{v_1}. \quad (22)$$

For DZ twins reared together, we have $M_1 = [(1-p) + p\gamma_1]M_0 - p\gamma_1 Z_0 + pX_1$ and $M'_1 = [(1-p) + p\gamma_1]M'_0 - p\gamma_1 Z_0 + pX'_1$, because these siblings share the same parents; biological considerations suggest that the correlation between M_0 and M'_0 is $\frac{1}{2}$. Therefore,

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

and the quantity $r^{MZT} - r^{DZT}$ can be readily obtained from Eqs. (22) and (23). 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 $M'_1 = [(1-p) + p\gamma_1]M_0 - p\gamma_1 Z'_0 + pX'_1$, because these siblings have the same genetic endowment, but different parents. If adoptive parents are independently drawn, the correlation of Z_0 and Z'_0 is zero, and so

$$r^{MZA} = [(1-p) + p\gamma_1]^2 \frac{v_0}{v_{1a}}. \quad (24)$$

Simple calculations using Eqs. (20) and (21) yield the expression for $r^{MZT} - r^{MZA}$ in Eq. (6).

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* parent²⁷ with time-(-1) mean Z_{-1} . By analogy with Eq. (5),

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

Z_{-1} has a similar expression, which we can use to substitute for Z_{-1} above. Iterating, we get

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

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 $2^{-t} v_0$. Also, M_0 is uncorrelated with observations; thus,

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

Finally, to compute v_{Z0} , note that, in steady state, it must be the case that $v_1 = v_{Z0}$; furthermore, clearly $\text{Cov}[M_0, Z_0] = r_0 \sqrt{v_0 v_{Z0}}$. Thus, we can substitute for the latter quantity in Eq. (20), assume that $v_1 = v_{Z0}$ and solve for v_{Z0} : we get

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

²⁷Adding a small fraction of adopted children in each generation has very limited quantitative effects.

References

- [1] Akabayashi, Hideo. (2006) “ An equilibrium model of child maltreatment” *Journal of Economic Dynamics and Control* Volume 30, Issue 6, June 2006, Pages 993-1025.
- [2] Bandura A. (1977): *Social Learning Theory*, Prentice-Hall.
- [3] Bandura A. and R.H. Walters (1963): *Social Learning and Personality Development*, New York: Holt, Rinehart & Winston.
- [4] 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.
- [5] Becker, G. S. (1981): *A Treatise on the Family*. Harvard University Press, Cambridge, Massachusetts.
- [6] 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.
- [7] Benabou, Roland and Jean Tirole (2003) “Intrinsic and Extrinsic Motivation” *Review of Economic Studies*, 70(3) (2003), 489-520.
- [8] Bernal, Raquel (2004): “Employment and Child Care Decisions of Mother’s and the Well-being of their Children,” manuscript, Northwestern University.
- [9] 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.
- [10] Bisin, A. and T. Verdier (2001): “Cultural Transmission, Socialization, and the Dynamics of Preferences,” *Journal of Economic Theory*, 97(1), 298-319.
- [11] 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.
- [12] Bjorklund, Anders, Jantti, Markus, and Gary Solon (2007): “Nature and Nurture in the Intergenerational Transmission of Socioeconomic Status,” NBER.
- [13] Black, Sandra, Devereux, Paul, and Salvanes, Kjell (2005): “The More the Merrier? The Effect of Family Size and Birth Order on Children’s Education,” *Quarterly Journal of Economics*, Vol. 120, N. 2 669-700.
- [14] 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.
- [15] Camerer, Colin F. (2003): *Behavioral Game Theory: Experiments in Strategic Interaction*; Princeton University Press.
- [16] Carneiro, Pedro, Heckman, James J. and Dimitriy V. Masterov (2002): “Labor Market Discrimination and Racial Differences in Premarket Factors,” manuscript, University of Chicago.

- [17] Chisholm, K. "A Three Year Follow-up of Attachment and Indiscriminate Friendliness in Children Adopted from Romanian Families." *Child Development* 69(4): 1092-1106.
- [18] 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.
- [19] Cunha, Flavio and James Heckman (2007): "The Technology of Skill Formation," NBER Working Paper 12840.
- [20] 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.
- [21] 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.
- [22] Ettinger, David and Philippe Jehiel (2006): "Towards a Theory of Deception," manuscript, PSE and University College London.
- [23] Fudenberg Drew and David K. Levine (1998): *Theory of Learning in Games*; MIT Press.
- [24] Goldberger, Arthur S. (1979), "Heritability," *Economica*, New Series, Vol. 46 n. 184, pp. 327-347.
- [25] 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.
- [26] Harlow, H. F., and Zimmerman, R. (1959). "Affectional responses in the infant monkey." *Science*, 130, 421-432.
- [27] Harris, Judith Rich (1995): "Where Is the Child's Environment? A Group Socialization Theory of Development," *Psychological Review*, Vol. 102, N. 3, 458-489.
- [28] Harris, Judith Rich (1998): *The Nurture Assumption*. The Free Press.
- [29] Herrnstein, Richard, and Charles Murray (1994) *The Bell Curve* The Free Press.
- [30] 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.
- [31] Pinker, S. (2002): *The Blank Slate: The Modern Denial of Human Nature*, Viking-Penguin.
- [32] 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.
- [33] Plomin, R., DeFries, J. C., McClearn, G. E., & McGuffin, P. (2001): *Behavioral Genetics*. (4th Edition). New York: Worth Publishers.
- [34] Reiss David, Neiderheiser Jenae M., Hetherington Mavis E. and Plomin Robert (2000): *The Relationship Code*. Harvard University Press.

- [35] Ridley, Matt, (2003): *Nature Via Nurture : Genes, Experience, and What Makes Us Human*, Harper Collins.
- [36] 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.
- [37] Sacerdote, Bruce (2002): "The Nature and Nurture of Economic Outcomes," *American Economic Review Papers and Proceedings*, vol. XCII p. 344–348.
- [38] 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.
- [39] Shonkoff, Jack, and Deborah Phillips, eds. 2000. *From Neurons to Neighborhoods: The Science of Early Childhood Development*. Washington, D.C.: National Academy Press.
- [40] Smith, Lones and Peter Sorensen (2006): "Informational Herding and Optimal Experimentation," manuscript, University of Michigan.
- [41] 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.
- [42] Turkheimer, Eric, Haley, Andreanna, Waldron, Mary, D'Onofrio, Brian, and Gottesman, Irving I. (2003): "Socioeconomic Status Modifies Heritability of IQ in Young Children," *Psychological Science* XIV 623-628
- [43] 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.
- [44] Vygotsky L.S. (1978). *Mind in Society*. Cambridge, MA: Harvard University Press.
- [45] 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.
- [46] Weinberg, Bruce. (2001) "An Incentive Model of the Effect of Parental Income on Children" *The Journal of Political Economy*, Vol. 109, No. 2. pp. 266-280.