Skill Formation and Parenting Styles: Evidence From a Structural Supermodular Game with Multiple Equilibria*

Marco Cosconati †

September 2012

Abstract

In this paper I develop a dynamic supermodular game of parent-child interaction with multiple equilibria. I allow households to select different equilibria, which I compute using a numerical algorithm developed by Echenique (2007). I use the model to assess the impact of parental incentives on the formation of skills, to characterize the equilibria which are more likely to generate the data and how they are related to family background characteristics.

Keywords: Games with strategic complementarities, skills formation

JEL Classification: C15, C18, C73.

*VERY PRELIMINARY: PLEASE DO NOT CITE OR QUOTE.

†Research Department, Bank of Italy, marco.cosconati@bancaditalia.it

1 Introduction to the parenting game

In order to place the parent-child game within the context of the child's development it is useful to distinguish two main phases. In the first phase, the *passive* phase, the child does not respond to parental inputs which mechanically determine his human capital and parents also determine his preferences which also evolve dynamically. In the second stage of life, the *active* phase, the child's preferences are more stable and the child maximizes his expected utility acting in a forward-looking fashion. In particular he reacts to parental inputs contributing to his own development. The game I will describe is suited to understand the role of parenting styles during the *active* phase, which typically coincide with adolescence. The willingness to pursue conflicting objectives generates the strategic interaction which is analyzed in this article.

General set-up. I analyze the strategic behavior of parents and children within a *dynamic stochastic game of complete information with strategic complementarities* [1]. There are two players, the *child* (he) and the *parent* (she). Time is discrete and the stage game is repeated T times. The state variables of the game are given by the child's skills which evolve over time according to a Markov process. Following the literature on skills formation [2]. I distinguish between cognitive (θ^c) and noncognitive skills (θ^{nc}).
All the parameters of the game as well as the child's skills are *common knowledge.* The initial condition of the game is θ_1, and it is taken as given and fully known by the players. *Players are rational and forward looking* and discount the future at rate δ[3]. At the beginning of each $t-$stage game, the parent selects a *parenting style* ps_t which is given by a pair (c, s): $c \in [0, 1]$ is the *share of parental income* y the child receives each period, $s \in [0, 1]$ is a measure of the degree of *supervision and support* parents give to their children in the knowledge acquisition process. *Simultaneously* the child selects an effort level a_t which, which is conceptualized in term of the time the child spends acquiring new knowledge, e.g. learning time. Because total available time is normalized to one leisure l_t equals $1 - a_t$. Given the skills of the child θ_t and the players' actions a new state θ_{t+1} realizes stochastically. The stage game is then repeated until the last period T.

Evolution of skills *New knowledge* -θ^c_{t+1}- is a function of the child's effort, of the current stock of human capital and of level of parental supervision s. The child's cognitive skills take J_c finite number of values, i.e. $\theta^c_t \in \langle \theta^c_1, \ldots, \theta^c_{J_c} \rangle$. I model the (stochastic) law of motion of the cognitive skills using an ordered logit framework:

$$\theta^c_{t+1} = \theta^c_j \Leftrightarrow q^c_{j-1} \leq \overline{\theta}^c_t \leq q^c_j \tag{1}$$

[1]See vives for an introduction.

[2]See Cunha and Heckman (2007), Cunha and Heckman (2008) and Cunha, Heckman, and Schennach (2010).

[3]It might be sensitive to assume that the child is more myopic than the parent. In practice I do not estimate the discount factor, which is not separately identified from the other parameters of the primitives. Therefore for ease of exposition, the parent and the child share the same δ.

where the q_j^c are critical cut-offs and the index $\overline{\theta}_t^c$ is given by

$$\overline{\theta}_t^c = \eta_1^c a_t + \eta_2 s_t + \eta_3^c \theta_t^c$$

The ordered logit model implies that

$$\Pr(\theta_{t+1}^c = \theta_j^c | \overline{\theta}_t^c) = \Big[\frac{\exp(q_j^c - \overline{\theta}_t^c)}{1 + \exp(q_j^c - \overline{\theta}_t^c)} - \frac{\exp(q_{j-1}^c - \overline{\theta}_t^c)}{1 + \exp(q_{j-1}^c - \overline{\theta}_t^c)}\Big] \tag{2}$$

It is easy to check that for any increasing function V

- $\mathbb{E}[V(\theta_{t+1}^c)|\overline{\theta}_t^c]$ is increasing in $(a_t, s_t, \theta_t^{nc})$
- $\mathbb{E}[V(\theta_{t+1}^c)|\overline{\theta}_t^c]$ has increasing differences for any pair of variables in $(a_t, s_t, \theta_t^{nc})$

which means that the distribution $F(\cdot|a, s, \theta^{nc})$ has increasing differences in any pair of variables. I conceptualize *noncognitve skills* in terms of the *stock of effort habits*, in short the child's *habits*. The law of motion of the habits is given by

$$\theta_t^{nc} = \theta_{t-1}^{nc} + \lambda(a_{t-1} - \theta_{t-1}^{nc}) \tag{3}$$

The parameter $\lambda \in [0, 1]$ indexes the speed with which leisure impacts habits; if $\lambda = 1$ habits in the current period collapse to the previous period effort time, while if $\lambda = 0$ the behavior of the child has not effect on future habits. In short θ_t^{nc} can be thought has the child's *discipline*, a measure of how much he is used to work. To better understand how parental incentives interact with the child's habits and his actions it is useful describe the child's preferences. I conceptualize *noncognitve skills* in terms of the *stock of effort habits*, in short the child's *habits*. To better understand how parental incentives interact with the child's habits and his actions it is useful describe the child's preferences. For future references it is useful to denote by $F^c(\theta_{t+1}^c|\theta_t^c, a_t, s_t)$ and by $F^{nc}(\theta_{t+1}^{nc}|\theta_t^{nc}, a_t)$ the law of motion of the cognitive/noncognitive skills respectively. The probability distribution of the skills is written in a compact way as $F(\theta_{t+1}|\theta_t, a_t, ps_t)$.

Children's preferences The *child* cares about his *leisure*, his consumption and his *adult human capital*. The payoff function of the child is given by:

$$v_t = \begin{cases} u(c_t, \theta_t^{nc}, a_t) & \text{if } t < T+1 \\ \Xi(\theta_{T+1}) & \text{if } t = T+1 \text{ (the game is over)} \end{cases} \tag{4}$$

I parametrize u as follows

$$u(c_t, \theta_t^{nc}, a_t) = -\exp\big[-\tau_0 c_t - \tau_1(1 - a_t + \gamma\theta_t^{nc})\big] \tag{5}$$

with τ_0, $\tau_1 > 0$ and $\gamma > 0$. As it is standard in the habit formation literature, the term $\gamma\theta_t^{nc}$ acts as a reference point in evaluating the disutility of a: higher θ^{nc} lower the negative impact of study time on the child's utility. Besides displaying constant absolute relative risk aversion in $c$and l, the utility function has the following properties

1. u is increasing in consumption and leisure (decreasing in a)
2. the marginal disutility of effort (utility of leisure) is a) decreasing in the habits and b) decreasing in consumption
3. the level of utility is decreasing in the habits

Property 1) is obvious and generates the trade-off the child faces between increasing his adult human capital and enjoying leisure time. Property 2b)-the payoff function has decreasing differences in consumption and leisure- is at the heart of the mechanism through which parental transfers generates a positive feedback on the child's current behavior. In short, it is *as if* the parent can "purchase" study time by transfering monetary resources to the child. Property 2a) allows for a propagation effect on future periods: higher the parameter γ more effective parental transfer are in letting the child to get used to work. Therefore there are both static and dynamic *strategic complementarities* between a and c. Property 3) is standard in habit formation models and it implies that if the child enjoyed a great deal of leisure in the past, learning is a particularly costly activity.

The final payoff of the child is a function of the skills cumulated at the end of the game. The valuation function Ξ can be thought as a proxy for future wages/lifetime income or, in the language of Cunha and Heckman, the child's "*adult human capital*". Following the CES example of Cunha and Heckman (2008) and Cunha et al. (2010):

$$\Xi(\theta) = \left[\tau_a(\theta^c)^\phi + (1-\tau_a)(\theta^{nc})^\phi\right]^{\frac{1}{\phi}} \tag{6}$$

with $\tau_a \in [0,1]$ and $\phi \in (-\infty, 1]$.

Parent's preferences The utility she enjoys at the end period t is given by the following formula:

$$w_t = \begin{cases} \omega y(1-c) - \beta s & \text{if } t < T+1 \\ \Pi(\theta_T) & \text{if } t = T+1 \text{ (the game is over)} \end{cases} \tag{7}$$

with β and ω strictly positive. According to this specification providing supervision to the child entails a psychic cost, while high transfers imply a higher opportunity cost in terms of private consumption.

To conserve on parameter the value parents attach to the child's adult human capital, is given by:

$$\Pi(\theta_T) = \iota\Xi(\theta_T) \tag{8}$$

where $\Xi(\theta_T)$ is written in (6) and $\iota > 0$.

To sum-up the effort of the child is increasing in his cognitive (noncognitive) skills through the positive interaction in F_c (u) between a and θ_c (θ^{nc}). Parental support has a indirect effect on future θ^c by increasing a's marginal productivity in F_c, and a direct effect given by it's own impact on future cognitive skills. Monetary transfers affect the child's cognitive skills through its impact on the effort which will then affect future noncognitive skills. Therefore I allow for an an endogenous cross-productivity of noncognitive skills on cognitive skills[4].

Interpretation and relationship to the literature. The *source of conflict* in the model between the parent and the child stems from a miscongruence in preferences. In particular the parent cares about the human capital of the child more than he does while holding a lower private valuation of the child's leisure.[5] This approach is consistent with the existing literature, see Weinberg (2001) and Akabyashi (2006). There are two main points of departure. The first regards the typical assumption that children are myopic, namely have a discount factor of zero. On the contrary the parent is forward looking and knows the child will eventually recognize the benefits of investing on her human capital. Although it capture the insight that parents "know better", a less extreme assumption- i.e. the child is more (not totally) myopic than the parent- seems more a realistic basis to model adolescent behavior. In this set-up the dynamic nature of the child's optimization problem allows the parent to credibly threaten of taking away privileges in case of poor performances. Obviously such a threat would not be effective if the child does not internalize the cost of future punishments.
The second difference regards the timing. These two papers model the game in terms of a principal-agent model[6] with the parent acting as leader, while I propose a simultaneous move game. Either way the researcher is making an assumption and there is really no compelling protocol to stick to. Finally for what concerns *link between the parent and the child*, I rely on strategic complementarities [7]. The role of my parenting style is closed in spirit to the "incentive schedule" of Akabyashi (2006) and allows the parent to influence the child's behavior even in absence of binding contracts, as in Weinberg (2001). The presence of dynamics has also implications for the optimal timing of parental incentives as in the literature on the technology of skills formation[8]. In particular the patterns of parenting across ages of the child -more supervision is in place at earlier ages- is consistent with a positive complementarity between effort and human capital: the parent stimulates effort which increases the child's future human capital and habits creating a virtuous circle. Having acquired a strong working habit the child will not need external incentives at later ages. This is why in this context the implications for the optimal timing of incentives is also a result of the dynamic behavior of the child, and not only of the features of skills production function as in the literature on

[4]The reverse type of cross productivity would be allowed if θ^c would affect the child's habits.

[5]In the words of Amy Chua:*"To get good at anything you have to work and children on their own never want to work, which is why it is crucial to override their preferences. This often requires fortitude on the part of parents because the child will resist."*

[6] In Weinberg (2001) commitment is possible while Akabyashi (2006) does not.

[7] A a linear specification of the specification of the child's flow utility $(1-a)\mu\epsilon$ enables to interpret parenting styles in terms of a proportional *composite* tax the parent imposes on the child's leisure. Such a tax is given by the combination of a subsidy (the allowances) and of "specific" taxes on leisure activities (the schedule of rules. I thank Lance Lochner for suggesting such an interpretation.

[8]See Cunha and Heckman (2007) and Cunha et al. (2010).

early intervention.

2 Markov Equilibria

In order to define the equilibrium it is useful to write the problem in sequential fashion. The history of the game is given by the sequences of actions and skills realizations $\mathcal{H}^t = (a^t, \theta^t, ps^t)$, where as in the standard notation the vector x^t is given by $x^t = (x_1, \dots, x_t)$. The problem solved by the child reads as follows:

$$\max_{a^T} \sum_{t=1}^{T} \mathbb{E}[v_t | ps^T, \mathcal{H}^t] \tag{9}$$

$$\text{s.t } \; a_t \in [0,1] \;\; \forall a_t \in a^T \tag{10}$$

where the expectation is taken over the shocks appearing into the skills productions functions. The problem solved by the parent is to find a sequence of parenting styles in order to:

$$\max_{ps^T} \sum_{t=1}^{T} \mathbb{E}[w_t | a^T, \mathcal{H}^t] \tag{11}$$

$$s.t. \;\; ps_t \in [0,1] \times [0,1] \;\; \forall ps_t \in ps^T$$

A pure strategy Nash equilibrium of the game is given by a sequence of effort levels and parental incentives $\langle a^T, m^T) \rangle_{t=1}^{T}$ such that i) a^T is a best response to the sequence μ^T after any history $\mathcal{H}^t$, and viceversa; ii) the strategies are a Nash equilibrium for every subgame subsequent to any history $\mathcal{H}^t$. The set $\mathcal{E}$ which contains all the Nash equilibria of the game may not be a singleton. The *game admits multiple equilibria* because it is one of simultaneous moves, therefore multiplicity arises as in any coordination game. Second, the objective functions of the players are not quasi-concave[9]. For tractability, as it is custom in dynamic stochastic games, I focus my attention to Markov equilibria. In particular I take θ_t -the child's skills at the beginning of the stage game- as a sufficient statistics for the history $\mathcal{H}^t$.

The relationships postulated so far can be summarized as follows

- The payoff of the players display increasing differences in the opponent's action

[9] Starting from the last period notice how, even assuming that the all the primitive functions are concave, there is no guarantee that the objective function of the parent is quasi-concave in (a, μ). This is the case even if the effort function $a(\mu, \theta)$ is concave in μ, because concavity is not an ordinal property. Relying on quasi-concavity is also problematic because the property is non closed under summation. Hence even if the future component of the payoff function is quasi-concave there is no guarantee that summing it to the flow utility delivers a quasi-concave function. The conditions provided by Quah and Strulovici (2010) to check the quasi-concavity did not prove useful. The non-concavity of the player's objective functions prevents the applicability of the first order approach in the principal-agent model. Cosconati (2011) provides sufficient conditions for the uniqueness of the equilibrium in the context of a simpler model using the CDFC condition.

- The payoff of the players display increasing differences in their own actions and the states
- The payoff of the players are increasing in the opponent's action

These properties, coupled with other technical properties such as the boundedness of the payoff functions and the stochastic nature of the transition function, are such that the existence of a MPE in ensured by Amir (2011)'s theorem[10].

2.1 Multiple equilibria across families

The estimation of games with multiple equilibria is a complicated matter, due to the well known indeterminacy (or incompleteness) problem, which makes the map from the structural parameters to the likelihood of observed events a correspondence rather than a function[11]. There are two prevalent approaches. In the context of discrete games, both static and dynamic, the approach developed by Hotz and Miller (1993) can be fruitfully applied to games, as done by Bajari, Benkard, and Levin (2007). The convenience of this approach relies on the fact that there is no need to solve for the equilibria of the game at any given parameter trial. Therefore the typical inner loop of the standard full-solution method, is not performed. In the context of static games, exploiting the particular structure of his game, Moro (2003) develops a two step estimator which also does not require to compute all the equilibria of the game [12]. The starting point of these papers is that computing all the equilibria can be undoable for large games. Both approaches hinge on the assumption that a unique equilibrium is played across markets/families. Bajari, Hong, and Ryan (2010) develop an algorithm which relaxes this assumption and allows the data generating process to be a mixture of likelihood functions, each associated to a specific equilibrium. They estimate a game of entry in auctions by computing all the pure and mixed strategy equilibria. This is achieved by using the software Gambit, developed by which solves systems of polynomial equations and inequalities. The authors provide conditions for the semi-parametric identification of the payoff function and of the selection's mechanism. Moreover they model and estimate the equilibrium selection rule. They allow the equilibria to be characterized by three "attributes". An equilibrium can be a pure or a mixed strategy equilibrium, Pareto optimal or inefficient and to maximize the sum of the utilities of the players or not. This approach, while allowing to understand how equilibria are selected, enables to simulate the model and run counterfactuals experiment without having to impose an equilibrium selection rule. My estimation procedure can be seen as a variant of their approach. I model (parametrically) the mechanism that maps those variables into the probability that an equilibrium is played.

Because the data on parent-child interaction are allowed to be generated by different equilibria I can use the model to address the following question: 1) what are the char-

[10] Curtat (1996) establishes existence for the case of infinite horizon. See also Vives (2007) for a characterization of the equilibria in dynamic games with complementarities.

[11] see Tamer (2003) for an explanation of the issue of incompleteness and Bajari, Hong, and Nekipelov (2010) for a survey of the methods in empirical IO.

[12] See Moro, Bisin, and Topa (2011) for an application to social interaction.

acteristics of the equilibria which are more likely to be played in American families? 2) How is the probability of selecting a given class of equilibria related to observable family characteristics (such as family income and race)?
The theory of supermodular games is extremely useful both in providing sensitive criteria to "classify" the equilibria over which the selection mechanism places a probability distribution, as well as in computing the set of equilibria. Roughly speaking finding Nash equilibria implies solving numerically a system of non-linear equations, typically equilibrium conditions, which admits more than one solution. From the computational point of view this can be a daunting task [13]. I approach the problem by implementing on a larger scale the algorithm developed by Echenique (2007), which I will describe later, showing how it can be applied to my setting.

2.2 Solution Method

Some Preliminaries

Before describing the algorithm it is useful to recall the following properties of GSC games. Consider a GSC with strategy space $S = [0,1] \times [0,1]$, call it Γ. The games has the following properties:

- Each player has a largest $\overline{\sigma}_i(s_{-i})$ and a smallest $\underline{\sigma}_i(s_{-i})$ best reply, for $i = \{c, p\}$, which are increasing in the strategy of the opponent (Topkis (1979) and Milgrom and Roberts (1990)) (increasing best responses).
- There exists a largest ($\underline{s}$) and a smallest equilibrium ($\overline{s}$) (ordering of equilibria).
- The equilibria can be Pareto ranked because the payoff of the parent is increasing in the effort of the child. Therefore $\overline{s}$ is the Pareto best equilibrium and $\underline{s}$ is the smallest one is the Pareto worste (welfare).
- Best-reply dynamics approach the interval $[\underline{s}, \overline{s}]$. Therefore starting from $(0,0)$ and $(1,1)$ best-reply dynamic converge to $\underline{s}$ and ($\overline{s}$), respectively. (Global stability of the extremal equilibria)
- The largest and the smallest equilibrium points increase with an in increase in θ (comparative statics of equilibria)
- Let $\Gamma(s_i)$ denote the game with an action space $s \geq s_i$. First, if Γ is GSC then $\Gamma(s_i)$ is a GSC. Moreover if s is a Nash equilibrium of Γ and $z \leq s$ then s is a Nash equilibrium of $\Gamma(z)$ (robustness to "truncation").

For the time being let's assume that the researcher is able to solve for the equilibria of any static GSC. Let N be the number of points in the state space Θ. Let $\langle V_T^c(\theta|e), V_T^p(\theta|e)\rangle$ denote the value functions of the child/parent in state θ associated to equilibrium $e \in \mathcal{E}_T(\theta)$, where $\mathcal{E}_T(\theta)$ is the equilibrium set in state θ and $\mathcal{E}_T(\Theta)$ the set of equilibria in all the states. Let's define the following

[13] As stated in Press, Teukolosky, Vetterling, and Flannery (2007) "*there are no good, general methods for solving systems of more than one nonlinear equations.". See McKelvey and McLennan (1996) for a survey on available algorithms.*

Definition 1. A sequence of equilibria $\langle e_t(\theta_i)\rangle_{i=1}^N$ is an *increasing sequence* or has the *increasing property* if the associated value functions $\langle V_t^c(\theta_i), V_t^p(\theta_i)\rangle_{i=1}^N$ are such that $V_t^c(\theta_i) \geq V_t^c(\theta_j)$ $V_t^p(\theta_i) \geq V_t^p(\theta_j)$ for any $\theta_i > \theta_j$.

Notice that the set of equilibria of the game with the increasing property, call it $\mathcal{E}^i(\Theta)$, is a subset of the all the pure strategy equilibria of the game. Therefore the increasing sequence criterion is effectively a refinement. The set $\mathcal{E}^i(\Theta)$ can be computed as follows:

- Proceeding backward, consider the T-period game. Compute, for every θ $\mathcal{E}_T(\theta)$ using Echenique's algorithm (decribed in section 2.2.1), which provides the equilibrium set $\mathcal{E}_T(\Theta)$.
  Construct $\mathcal{E}_T^i(\Theta)$. Notice that $\mathcal{E}_T^i(\Theta)$ is non empty because the highest equilibrium in $\mathcal{E}_T(\Theta)$ is Pareto optimal and increasing in θ (the game has *positive externalities*, i.e. the payoff functions of the players is increasing in the opponent's action). The number of continuation values I consider in the next step equals the cardinality of the set $\mathcal{E}_T^i(\Theta)$
- At $T-1$ at a given state θ, for all the possible continuation values induced by $\mathcal{E}_T^i(\Theta)$, use Echenique's algorithm to compute $e_{T-1}(\theta)$. Refine $\mathcal{E}_{T-1}(\Theta)$ to obtain $\mathcal{E}_{T-1}^i(\Theta)$
- Keep iterating this procedure for periods $T-2, T-2, \ldots, 1$

After applying the admissible sequences refinement the total number of equilibria equals the number of elements in $\mathcal{E}_1^i(\Theta)$. It is legitimate to wonder on which basis such a refinement can be justified. Roughly speaking the *increasing sequence* criterion (IS) is a way of "supermodularizing" the game. As argued by Echenique (2003) static complementarities not necessarily translate into dynamic complemementarities. Restricting the attention to equilibria in which the continuation value is increasing in the states is akin to construct sequences of supemodular payoffs in which the future components (the value functions) are indeed supermodular.
Further, given the existence of positive externalities, the IS requirement can be interpreted as requiring the equilibria to be non decreasing in the state. In other words if, counterfactually with respect to the timing of the game, the amount of skills exogenously increased *at a given time period*, it has to be that the child increases his effort and the parent selects a "higher" parenting style, e.g. a parenting style which increases the parent's expected utility. In this sense the Pareto-increasing property of the IS criterion is a way of ordering the choice set of the parent. For instance if the parent's choice set was an interval of the real line, the IS criterion would be equivalent to ask that Samuelson's correspondence principle (CP) holds. Because Echenique (2002) shows that (CP) holds in GSC, the AP criterion is also a way making (CP) operational.

2.2.1 Echenique's algorithm

The algorithm to compute the pure strategy Nash equilibria is based on the properties of the equilibria of supermodular games. Mixed equilibria are not considered on the ground of some theoretical considerations [14]. For ease of exposition consider the T-period stage game. Fix a grid on the action spaces of the player of n points, so that $\langle a_i \rangle_{i=1}^n$ and $\langle ps_i \rangle_{i=1}^n$ are the discretized versions of the actual action spaces of the child and of the parent, respectively. In particular let $\pi^c(\theta, s)$ be the payoff function of the child in a state θ conditional on $e = (a, ps)$ and $\pi^p(\theta, s)$ be defined analogously. Let $\beta_{c,\Gamma}(\theta, ps) = \arg\max_a \pi^c(\theta, s)$ and let $\beta_{p,\Gamma}(\theta, a) = \arg\max_p s\pi^p(\theta, s)$ and let $\beta_\Gamma(s, \theta) = \beta_{c,\Gamma}(\theta, s) \times \beta_{p,\Gamma}(\theta, s)$ be the game's best response correspondence. Let $\mathcal{E}(\theta)$ be the set of Nash equilibria of the game. The algorithm consists of the following steps:

step 1 Calculate $\inf \beta_{i,\Gamma}(\theta, e)$ and $\sup \beta_{i,\Gamma}(\theta, e)$, for $i = c, p$, that is the "lowest" and "highest" optimal response of the players for each possible action (on the grid) taken by the opponent. Approximate these best responses by interpolating the functions I calculated at the grid points. Find the extremal equilibria by iterating the best response correspondence. This is achieved by using the Robinson-Topkis algorithm[15]:

- $\underline{T}(\theta, e)$: starting from an initial guess $s_0 = s$, the sequence $\langle s^k(\theta) \rangle$ with $s^{k+1} = \inf \beta_\Gamma(e^k)$ converges to the lowest equilibrium.
- $\overline{T}(\theta, e)$: Given an initial guess s_0 the sequence obtained from $e^{k+1} = \sup \beta_\Gamma(e^k)$ converges to the highest equilibrium.

Call the extremal equilibria $\underline{e}_1 = (\underline{a}_1, \underline{ps}_1)$ and $\overline{e}_1 = (\overline{a}_1, \overline{ps}_1)$

step 2 Consider the game $\Gamma(\theta, (\underline{a}_1 + 1, \underline{ps}_1))$, where $\underline{a}_1 + 1$ is the point on the grid right next to $\underline{a}_1$. Use $\underline{T}(e, \theta)$ to find the smallest equilibrium of this restricted game, call it $e^1 = (a_1, ps_1)$. At this point it needs to be checked that e_1 is indeed an equilibrium (Echenique calls this the "check phase"). This is accomplished by simply checking that there is no profitable deviation for the child in the interval $[\underline{a}_1, \underline{a}_1 + 1]$. Why is it enough to consider this interval? Because s_1 is an equilibrium of the game $\Gamma(\theta, (\underline{a}_1 + 1, \underline{ps}_1))$ strategies $a > a_1$ cannot constitute a profitable deviation; on the other hand $\beta_{c,\Gamma}(\theta, ps_1) \geq \beta_{c,\Gamma}(\theta, \underline{ps}_2) = \underline{a}_1$, only strategies greater than $\underline{a}_1$ need to be considered. If s^1 passes this check, it is added to the set $\mathcal{E}$. Analogously consider the game $\Gamma(\theta, (\underline{a}_1, \underline{ps}_1 + 1))$ and apply $\underline{T}(e, \theta)$ to this game to find a new equilibrium (if there is one). Call it $e^2 = (a_2, ps_2)$.

step 3 Apply steps 2-3 to the games $\Gamma(\theta, (a_1 + 1, ps_1), \Gamma(\theta, (a_1, ps_1 + 1), \Gamma(\theta, (a_2 + 1, ps_2)$ and $\Gamma(\theta, (a_2, ps_2 + 1)$

step 4 Keep implementing steps 2-3 for each Nash equilibrium e^k found, until $e^k = \overline{s}$.

Echenique (2007) shows that the algorithm computes all the pure Nash equilibria of a GSC with finite strategies and that the algorithm is generally fast.

Efficient Application of the Algorithm GSC games are characterized by the following property: the smallest and largest equilibrium, $\underline{s}(\theta)$ and $\overline{s}(\theta)$ are increasing in θ.

[14] See Vives (1990) and Echenique and Edlin (2004).
[15] The proof that iterating the best responses, is provided by Topkis (1998) and Robinson (1951).

This characterization suggests an efficient way of calculating the equilibria through Echenique's algorithm. Given a discretization of the state space $\Theta = ((\underline{\theta}^c, \underline{\theta}^{nc}), \dots, (\underline{\theta}_i^c, \underline{\theta}_j^{nc}), (\overline{\theta}^c, \overline{\theta}^{nc}))$, with N grid points for both type of skills perform the following θ^{nc}-loop:

- iteration 1 Consider the smallest state $\underline{\theta} = (\underline{\theta}^c, \underline{\theta}^{nc})$ and apply $\underline{T}(e, \underline{\theta})$ with an initial guess $e_0 = (0, 0)$ and compute the set of equilibria $\mathcal{E}(\underline{\theta})$
- iteration 2 Consider the state $\theta^2 = (\underline{\theta}^c, \underline{\theta}^{nc} + 1)$ and use as initial guess e_0 for the operator $\underline{T}(e, \theta^2)$ the point $\underline{e}(\underline{\theta})$. Let $\underline{e}(\theta^2$ be the lowest equilibrium of the game.

$\vdots$

- iteration j Consider the state $\theta^{j-1} = (\underline{\theta}^c, \underline{\theta}^{nc} + j - 1)$ and use as initial guess e_0 for the operator $\underline{T}(e, \theta^{j-1})$ the point $\underline{e}(\theta^{j-1})$

$\vdots$

- iteration N Consider the state $\theta^{N-1} = (\underline{\theta}^c, \underline{\theta}^{nc} + N - 1)$ and use as initial guess e_0 for the operator $\underline{T}(e, \theta^{N-1})$ the point $\underline{e}(\theta^{N-1})$

Therefore, exploiting the comparative statics of the equilibrium set of GSC, the algorithm provides more accurate guesses for generating the sequence which converges to the smallest equilibrium of each game $\Gamma(\theta)$. As θ^{nc} increases the operator $\underline{T}(e, \theta)$ will (generically) need less iterations to find $\underline{e}$. To compute the set $\mathcal{E}(\theta)$ I perform an analogous θ^c-loop.

2.3 Simulating the model

There are several advantages of adopting the approach of Bajari et al. (2010), e.g. including the equilibrium selection as part of the econometric model. First, it allows to gain a better understanding how equilibria are selected. For instance, in my application one could compute, given the estimated parameters, which fraction of the equilibria is the Pareto optimal. Secondly the inclusion of λ allows to simulate the model and perform counterfactuals. They specify an equilibrium selection probability which allows to randomly draw an equilibrium which induces a probability distribution on the observed outcomes. Taking many draws from the selection probability function one can simulate the model multiple times and construct a simulated method of moments estimator. In here however one has to take into account that not all the set of equilibria are taken into account. Namely I consider only equilibria in which satisfy the IS criterion. As a result the selection rule needs to be adjusted to take into account this initial refinement.
An intuitive way of taking advantage of the ranking among equilibria based on Pareto optimality suggests the following procedure to pin down an equilibrium, consistently with the IS restriction needed to compute the equilibria. Once an equilibrium has been selected the model can be simulated many times and the parameters estimated by SML or MSM. Let $(\theta_1, \dots, \theta_N)$ denote the points of the state space, ordered from the lowest to the highest. Let $\#N_T(\theta_j)$ with $j = 1, \dots, N$ denote the number of equilibria in state θ_j at time T and let $\mathcal{E}_T(\theta_j) = (e_T^1(\theta_j), \dots, e_T^{\#N_T(\theta_j)}(\theta_j))$ be ordered according to their Pareto

ranking. An orderded logit/probit provides a natural equilibrium selection criterion. Let:

$$d = \beta X + \epsilon \tag{12}$$

denote the index which determines which equilibrium is played together with a set of cut-offs which I denote by c_j with $j = 1, \ldots \# N_T(\theta_j) - 1$ to be estimated along with the other parameters. Let ϵ be distributed according to a logistic/normal distribution. The β's are themselves parameters of interest to be estimated with the other structural parameters. A positive β implies that a given family characteristics, say parental education, increases the chances that an efficient outcome is reached. At time T the probability that equilibrium j is played in state θ_1 is given by

$$\lambda(e_T^j(\theta_1), \beta, X) = \Pr(c_{j-1} \le d \le c_j) \quad \text{with} \ \ j \in (1, \ldots, \# N_T(\theta_1))$$

Consider then the value functions $\langle V_T^c(e_T^j(\theta_1)), V_T^p(e_T^j(\theta_1)\rangle$ associated to the equilibrium selected in state θ_1, $e_T(\theta_1)$. Let

$$n(\theta_2|e_T(\theta_1)) \in (1, \ldots, \# N_T(\theta_2)) : V_T^k(e_T^j(\theta_2)) \ge V_T^k(e_T(\theta_1)) \ \text{ for any } \ \ j \ge n_T(\theta_2|e_T(\theta_1)); \ k = c, p$$

In words $n_T(\theta_2|e_T(\theta_1))$ is such that, *given the equilibrium selected in state* θ_1, any equilibrium $e_T^j(\theta_2)$ with an index j *lower* than $n_T(\theta_2|e_T(\theta_1))$ does *not* satisfy the IS criterion. To pin down an equilibrium at θ_2 which satisfies the IS criterion I can take another draw from the restricted set $(n_T(\theta_2|e_T^j(\theta_1)), \ldots, \# N_T(\theta_2))$ using the truncated distribution

$$\lambda^r(e_T^j(\theta_2), \beta, X|e_T(\theta_1)) = \frac{\lambda(e_T^j(\theta_2), \beta, X)}{\sum_{k=n_T(\theta_2)}^{\# N_T(\theta_2))} \lambda(e_T^k(\theta_2), \beta, X)} \quad \text{with} \ \ j = n_T(\theta_2|e_T(\theta_1)), \ldots, \# N_T(\theta_2)$$

where

$$\lambda(e_T^j(\theta_2), \beta, X) = \Pr(c_{j-1} \le q_2 \le c_j)$$

is calculate using the index d_2 in (12) and the distribution of ϵ. The full steps to draw an increasing equilibrium are:

- At time T
  1. Start from θ_1. Draw an equilibrium $e_T(\theta_1)$
  2. Move to θ_2 and compute $n_T(\theta_2|e_T(\theta_1))$. Draw an equilibrium $e_T(\theta_2)$ using the induced λ^r
  3. Keep iterating step 2 until the state space has been exhausted
  4. Store in memory the admissible sequence $\mathbf{V}_T(\Theta)$ given by the draws of the equilibria obtained by performing steps 1-3
- Move to $T - 1$
  1. Consider θ_1 and given the continuation values in $\mathbf{V}_T(\Theta)$ draw an equilibrium from the distribution $\lambda(e_{T-1}^j(\theta_1), \beta, X)$ in the support $(1, \ldots, \# N_{T-1}(\theta_1))$

2. Move to θ_2 and compute $n_{T-1}(\theta_2|e_{T-1}(\theta_1))$. Draw an equilibrium $e_{T-1}(\theta_2)$ using the induced λ^r
3. Keep iterating step 3 until the state space has been exhausted
4. Store in memory the admissible sequence $\mathbf{V}_{T-1}(\Theta)$ given by the draws of the equilibria obtained by performing steps 1-3

- Move to $T-2$ and performs analogous steps 1-4.
- Proceed backwards until period 1 has been reached.

2.3.1 Constructing the likelihood function

Let $\overline{\Theta}$ denote the set of parameters to be estimated through maximum likelihood, which consists of structural parameters and the β's parameters entering the equilibrium selection mechanism. The likelihood function is a mixture of likelihood functions, each associated to a given equilibrium and weighted by the probability that a given equilibrium is selected. Notice that because at any given equilibrium, there is no unobserved state variable such as random shock to preferences, the map between the state θ and the actions of the players is deterministic. Therefore the likelihood function associated to each equilibrium is degenerate. To deal with this issue I assume that all the variables are measured with error. I measure the effort exerted by the child using self-declared time spent doing homework. I adopt an hurdle model to deal with the cases in which the child declares not to study, e.g. the zeros. Let a^o denote an observed value of the child's effort. The hurdle model requires to specify the following conditional probability:

$$\Pr(a^o = 0|a) = \rho_a = \frac{1}{1 + \exp(\gamma_a a)}$$

as well as a conditional probability of observing a strictly positive measure of effort. I adopt a truncated normal distribution $g(\cdot|a)$ with mean $\lambda_{0,a} + \lambda_{1,a} a$ and variance ξ_a. The measurement error model is written as follows:

$$\mathbf{g}(e^o) = \begin{cases} \rho_a & \text{if } a^o = 0 \\ \frac{1-\rho_a}{1-G(0|a)} g(e^o|a) & \text{if } a^o > 0 \end{cases}$$

where G denotes the CDF of the density $g(\cdot|a)$.
An analog hurdle model is used to reconcile observed and endogenous parental transfers, which are measured by the self-declared amount of allowances the child receives. Parental support s is proxied by a number of variables obtained from the responses to questions made in the Autonomy/Parental Control Section of the NLSY97 questionnaire. The questions include how often the parent does each of the following:

- Praises the child for doing well
- Criticizes the child or the child's ideas
- Helps the child to do things important for the child

- Blames the child for his/her ideas
- Makes plans with the respondent and cancels for no good reasons.

DEAL with the issue of discrete measures.

For what concerns the measurement of the cognitive skills I follow Cunha et al. (2010) in specifying a linear factor model which allows to use multiple measures:

$$Y_j = m^j_{0,c} + m^j_{1,c}\theta^c + \varepsilon_c \quad \text{for} \quad j = 1, \dots, M \tag{13}$$

where M is the number of available measurements for cognitive skills. the m's are the factor loading of measurement j; ε_c is a measurement error term distributed according to a normal with zero mean and variance σ^2_c [16]. The measures I use to proxy for θ^c are the PIAT math test scores and the GPA achieved by the child at the end of the academic year.
The likelihood function is then simulated by drawing an equilibrium from the set of all increasing equilibria $\mathcal{E}_1(\Theta)$, calculating the associated likelihood using the measurement error and averaging over draws. **EXACT EXPRESSION TO BE WRITTEN**

[16] Cunha et al. (2010) provide conditions for the identification of the parameters in (13) as well as of the distribution of the θ's and of the ε_c.

References

Aguirregabiria, V. and P. Mira (2007, January). Swapping the nested fixed point algorithm: a class of estimators for discrete markov decision models. *Econometrica 75*(1), 1–54.

Akabayashi, H. (1995). A dynamic equilibrium model of parenting and child development with changing preferences and beliefs.

Akabyashi, H. (2006). An equilibrium model of child maltreatment. *Journal of Economic Dynamics and Control 30*(6), 993–1025.

Amir, R. (2011, September). Discounted supermodular stochastic games: Theory and applications.

Bajari, P., L. Benkard, and J. Levin (2007, September). Estimating dynamic models of imperfect competition. *Econometrica 75*(5), 1331–1370.

Bajari, P., H. Hong, and D. Nekipelov (2010, December). Game theory and econometrics: a survey of some recent research.

Bajari, P., H. Hong, and S. P. Ryan (2010, September). Identification and estimation of a discrete game of identification and estimation of a discrete game of complete information. *Econometrica 78*(5), 1529–1568.

Baumrind, D. (1968, Fall). Authoritarian vs. authoritative parental control. *Adolescence* (3), 255–272.

Becker, G. (1974, November). A theory of social interactions. *Journal of Political Economy 82*, 1063–1094.

Becker, G. S. and C. Mulligan (1997, August). The endogenous determination of time preferences. *The Quarterly Journal of Economics 112*(3), 720–758.

Benabou, R. and J. Tirole (2002, August). Self-confidence and personal motivation. *Quarterly Journal of Economics 117*(3), 871–915.

Bergstrom, T. (1989, October). A fresh look at the rotten kid theorem–and other household mysteries. *Journal of Political Economy 97*(5), 1138–1159.

Bernstein, N. I. (2001). *How to Keep Your Teenager Out of Trouble and What to Do if You Can't.* Workman.

Bornstein, M. H. (Ed.) (1991). *Cultural approaches to parenting.* L. Erlbaum Associates.

Bursztyn, L. and L. Coffman (2011, March). The schooling decision: Family preferences, intergenerational conflict, and moral hazard in the brazilian favelas.

Case, A. C. and L. F. Katz (1991, May). The company you keep: the effects of family and neighborhood on disadvantaged youths.

Chiappori, P.-A. and I. Werning (2002, April). Comment on “rotten kids, purity and perfection”. *Journal of Political Economy 110*(2), 457–480.

Cornes, R. and E. Silva (1999, April). Rotten kids, purity and perfection. *Journal of Political Economy 107*(5), 457–480.

Cosconati, M. (2009). *Parenting Style and the Development of human capital in children*. Ph. D. thesis, University of Pennsylvania.

Cosconati, M. (2011). The impact of birth order and family characteristics on parenting practices: Theory and evidence.

Cunha, F. and J. Heckman (2006). Formulating, identifying, and estimating the technology for the formation of skills, mimeo.

Cunha, F. and J. J. Heckman (2007). The technology of skill formation. *American Economic Review, Papers and Proceedings 97*(2), 31–47.

Cunha, F. and J. J. Heckman (2008, Fall). Formulating, identifying, and estimating the technology for the formation of skills. *Journal of Human Resources 43*(4), 738–782.

Cunha, F., J. J. Heckman, L. J. Lochner, and D. Masterov (2006). *Handbook of Economics of Education*. Eric Hanushek and Frank Welch. Amsterdam: North Holland.

Cunha, F., J. J. Heckman, and S. M. Schennach (2010, May). Estimating the technology of cognitive and non cognitve skill formation. *Econometrica 3*, 883–931.

Curtat, L. (1996). Markov equilibria of stochastic games with complementarities. *Games and Economic Behavior 17*, 177–199.

de Fraja, G., T. Oliveira, and L. Zanchi (2010). Must try harder. evaluating the role of effort on examination results. *Review of Economics and Statistics 92*, 577–597.

Dornbusch, S. M., P. L. Ritter, P. H. Leiderman, D. F.Roberts, and M. J. Fraleigh (2007, October). The relation of parenting style to adolescent school performance. *Child Development 58*, 1244–1257.

Echenique, F. (2002, March). Comparative statics by adaptive dynamics and the correspondence principle. *Econometrica 70*(2), 833–844.

Echenique, F. (2003). Extensive-form games and strategic complementarities. *Games and Economic Behavior 46*(2), 348–364.

Echenique, F. (2007). Finding all the equilibria in games of strategic complements. *Journal of Economic Theory 135*, 514–532.

Echenique, F. and A. S. Edlin (2004, September). Mixed strategy equilibria in games of strategic complements are unstable. *Journal of Economic Theory 118*(1), 61–79.

Eckstein, Z. and K. Wolpin (1999, November). Why youths drop out of high school: the impact of preferences, opportunities, and abilities. *Econometrica 67*(6), 1295–1339.

Edlin, A. S. and C. Shannon (1998). Strict monotonicity in comparative statics. *Journal of Economic Theory 81*, 201–219.

Fershtman, C. and A. Pakes (2011, March). Dynamic games with asymmetric information: A framework for empirical work.

Fiorini, M. and M. Keane (2011, February). How the allocation of children's time affects cognitive and non-cognitive development.

Grieco, P. (2010, November). Discrete game with flexible information structures: An application to local grocery markets.

Heckman, J. J., J. Stixrud, and S. Urzua (2006, July). The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior. *Journal of Labor Economics 24*(3), 411:482.

Hotz, J., L. Hao, and G. Z. Jin (2008, April). Games parents and adolescents play: Risky behaviors, parental reputation, and strategic transfers. *Economic Journal 118*(528), 515–555.

Hotz, J. and R. A. Miller (1993). Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies 60*(3), 497–529.

Lizzeri, A. and M. Siniscalchi (2008). Parental guidance and supervised learning. *Quarterly Journal of Economics, forthcoming*.

McKelvey, R. and A. McLennan (1996). The computation of equilibrium in finite games. In H. Amman, D. A. Kendricks, and J. Rust (Eds.), *Handbook of Computational Economics*, pp. 87–142. Amesterdam:Elsevier.

Milgrom, P. and J. Roberts (1990). Rationazability, learning, and equilibrium in games with strategic complementarities. *Econometrica 58*(6), 1255–1277.

Milgrom, P. and C. Shannon (1994, January). Monotone comparative statics. *Econometrica 62*(1), 157–180.

Moro, A. (2003, April). The effect of statistical discrimination on black-white wage differentials: Estimating a model with multiple equilibria. *International Economic Review 44*(1), 467–500.

Moro, A., A. Bisin, and G. Topa (2011, January). The empirical content of models with multiple equilibria in economies with social interaction.

Pabilonia, S. W. (2001, Autumn). Evidence on youth employment, earnings, and parental transfers in the national longitudinal survey of youth 1997. *Journal of Human Resources 36*(4), 795–882.

Press, W. H., S. A. Teukolosky, W. T. Vetterling, and B. P. Flannery (2007). *Numerical Recipes*. Cambridge University Press.

Quah, J. and B. Strulovici (2010, June). Aggregating the single crossing property: Theory and applications to comparative statics and bayesian games.

Robinson, J. (1951). An iterative method of solving a game. *The Annals of Mathematics 54*(2), 296–301.

Stinebrickner, T. R. and R. Stinebrickner (2008). The causal efect of studying on academic performance. *The B.E. Journal of Economic Analysis and Policy 8*, 1–53.

Sweeting, A. (2008). The strategic timing of radio commercials: An empirical analysis using multiple equilibria, *mimeo*.

Tamer, E. (2003, January). Incomplete simultaneous discrete response model with multiple equilibria. *Review of Economic Studies 70*(1), 147–165.

Todd, P. and K. Wolpin (2003, February). On the specification and estimating of the production production function for cognitive achievement. *Economic Journal 113*(485), F3–F33.

Todd, P. and K. Wolpin (2007, Winter). The production of cognitive achievement in children: Home, school and racial test score gaps. *Journal of Human Capital 1*(1).

Topkis, D. M. (1978). Minimizing a submodular function on a lattice. *Operation Research 26*, 305–321.

Topkis, D. M. (1979). Equilibrium points in nonzero-sum-n-person submodular games. *SIAM Journal of Control and Optimization 17*(6), 773–787.

Topkis, D. M. (1998). *Supermodularity and Complementarity*. Princeton University Press.

Van Zandt, T. (2002). An introduction to monotone comparative statics.

Vives, X. (1990). Nash equilibrium with strategic complementarities. *Journal of Mathematical Economics 19*(3), 305–321.

Vives, X. (2005). Games with strategic complementarities: New applications to industrial organization. *International Journal of Industrial Organization 23*, 625–637.

Vives, X. (2007, February). Strategic complementarities in multi-stage games.

Weinberg, B. (2001, April). An incentive model of the effect of parental income on children. *Journal of Political Economy 109*(2), 266–289.

Zandt, T. V. and X. Vives (2007). Monotone equilibria in bayesian games of strategic complemetarities. *Journal of Economic Theory 134*, 339–360.