UNIVERSITY OF CAMBRIDGE
Faculty of Economics

Institute for
New Economic Thinking

CAMBRIDGE WORKING PAPERS IN ECONOMICS
CAMBRIDGE-INET WORKING PAPERS

Parenting types

Christopher Rauh
University of Cambridge

Laëtitia Renée
University of Cambridge

Abstract

In this paper we measure parenting behavior through unsupervised machine learning in a panel following children from age 5 to 29 months. The algorithm classifies parents into two distinct behavioral types: "active" and "laissez-faire". Parents of the active type tend to respond to their children's expressions and describe to children features of their environment, while parents of the laissez-faire type are less likely to engage with their children. We find that parents' types are persistent over time and are systematically related to socio-economic characteristics. More-over, children of active parents see their human capital improve relative to children of parents of the laissez-faire type.

Reference Details

2110 Cambridge Working Papers in Economics
2021/06 Cambridge-INET Working Paper Series

Published 22 January 2021

Key Words Parenting styles, human capital, latent Dirichlet allocation, inequality, machine learning

Websites www.econ.cam.ac.uk/cwpe
www.inet.econ.cam.ac.uk/working-papers

Parenting types

By Christopher Rauh and Laëtitia Renée*

Draft: November 17, 2020

In this paper we measure parenting behavior through unsupervised machine learning in a panel following children from age 5 to 29 months. The algorithm classifies parents into two distinct behavioral types: "active" and "laissez-faire". Parents of the active type tend to respond to their children's expressions and describe to children features of their environment, while parents of the laissez-faire type are less likely to engage with their children. We find that parents' types are persistent over time and are systematically related to socio-economic characteristics. Moreover, children of active parents see their human capital improve relative to children of parents of the laissez-faire type.

Keywords: Parenting styles; human capital; latent Dirichlet allocation; inequality; machine learning

1. Introduction

Early childhood investments have been shown to be crucial for children's human capital development (Cunha, Heckman and Schennach 2010, Del Boca, Flinn and Wiswall 2014, Attanasio, Meghir and Nix 2020). Parental time investments generally are captured through different activities parents engage in with their children, such as visits to museums or the frequency of a parent reading to their children. The number of activities considered is either

* Rauh: University of Cambridge, Trinity College Cambridge (email: cr542@cam.ac.uk). Renée: McGill (email: laetitia.renee@mail.mcgill.ca). Rauh would like to thank the FRQSC for financial support (grant number 2020-NP-267422).

vast or restricted arbitrarily. When many investments are considered, they generally are combined (log-)linearly in latent factor models.[1] More recently the debate about parenting styles has emerged (Doepke and Zilibotti 2017), which discusses how variations in economic conditions predict parenting styles in terms of altruism and paternalism over time and across space.

This paper develops a new methodology to measure parenting types using unsupervised machine learning. The advantage of this approach is that it allows aggregating any number of granular parental activities in a non-linear fashion. Moreover, the resulting parenting types are interpretable. When restricting parents to two types, we find that parents can be classified into "active" and "laissez-faire". Active parents are more likely to be supportive of their children's progress and speak directly to their child, while laissez-faire parents are characterized by hardly interacting with their children in the presence of the interviewer.

We contribute to three strands of literature. First, we define a new way of dealing with the large dimensionality and complexity of parental activities in order to understand how parental investments impact children's human capital accumulation. We show that the uncovered parenting types are predictive of future human capital above and beyond the predictive power parental socio-economic characteristics or child fixed effects.[2]

Second, we contribute to the literature concerned about parenting styles (Cunha 2015, Doepke and Zilibotti 2019, Doepke, Sorrenti and Zilibotti 2019, Cobb-Clark, Salamanca and Zhu 2019, Agostinelli et al. 2020).[3] They generally draw the distinction of parenting styles between permissive, authoritarian, or authoritative, the choice of which depends on parental levels of altruism and paternalism and environmental factors such as returns to skills or taxation. The empirical approaches tend to classify parenting styles based on a single binary response to a survey question, such as how important obedience is for a respondent. Our ap-

[1]Parenting is characterized by a complex set of interactions and decisions. See Draca and Schwarz (2018) for a discussion on why linear combinations of features with the highest degrees of variance in the data may not provide optimal summaries of complex data generating processes.

[2]Despite the intuitive results we cannot claim causal effects due to the lack of an exogenous shock to parenting styles.

[3]Del Boca et al. (2019) propose a model in which parental types are not merely the outcome of utility maximization by the parents but the result of a bargaining process with the children.

proach allows capturing parenting styles based on many questions with complex interactions. Advantages of our data on parental activities is that they are not self-reported, but are observed and recorded by the enumerator, which should help to reduce systematic measurement error, and are the same set of actions observed across multiple survey waves.

Third, we add to the rapidly growing use of machine learning in Economics to classify behavioral types. The latent Dirichlet allocation (LDA) was originally developed by computer scientists Blei, Ng and Jordan (2003). The underlying idea is to classify text documents into a mixture of small number of topics. One key is that the topics are not predefined but are backed out through co-occurrence. We apply the same idea of topics to behavioral types. Other approaches to classifying behavioral types using LDA are Bandiera et al. (2017) who classify CEOs using detailed time-use surveys and find that CEOs distinct behavior affects firm performance. Draca and Schwarz (2018) use LDA to measure political ideology. We contribute to this literature by using LDA to classify investment behavior by parents and look at its relation to human capital accumulation in very early childhood.

2. Data

We use the Québec Longitudinal Study of Child Development (QLSCD), a detailed panel of a representative sample of families from Québec, a province in Canada, with a baby born between October 1997 and July 1998. More specifically, we focus our work on the 1,985 families who participated in the first three waves of the panel, conducted when the designated baby was 5, 17 and 29 months old.

We rely on the Observations of Family Life (OFL) instrument filled by the enumerator at the end of the annual interview. It includes observations made during the interview about the behaviour of the key respondent –the mother in 99% of the cases– and her interactions with her baby. This has the advantage of not relying on self-reported behavior which is common in the human capital literature and a potential source of bias.

We exclude mother-children pairs for whom the OFL instrument was not completed at child ages 5, 17 or 29 months because the child was sleeping. We end-up with a sample of 1,443 mother-children pairs. Table 1 describes the socio-economic characteristics of the families.

We focus our analysis on the ten variables from the OFL instrument that assess the behavior of the interviewed mother toward her child. Table 2 displays descriptive statistics for these variables. We see that some parental actions are highly dependent on the age of the child. For instance, the share of parents regularly checking on their child decreases from 72% when the child is 5 months old to 32% when the child is 29 months old.

Table 1—: Descriptive statistics

	N	Prop.
Number of siblings		
No sibling	656	45.5
One	560	38.8
Two or more	227	15.7
Household type		
Two-parent	1,179	81.7
Blended	159	11.0
Single-parent	102	7.1
Missing	3	0.2
Mother's age		
Less than 25	332	23.0
25-29	446	30.9
30-34	469	32.5
35 and more	195	13.5
Missing	1	0.1
Mother born outside Canada		
No	1,301	90.2
Yes	140	9.7
Missing	2	0.1
Language spoken at home		
French	1,176	81.5
Other	265	18.4
Missing	2	0.1
Mother education		
High school degree or less	380	26.3
Some college education	681	47.2
College degree	380	26.3
Missing	2	0.1
Parental working status		
Two-parents: both work	984	68.2
Two-parents: one works	304	21.1
Two-parents: none work	45	3.1
Single-parent: works	38	2.6
Single-parent: does not work	58	4.0
Missing	14	1.0
Below poverty threshold		
No	1,106	76.6
Yes	317	22.0
Missing	20	1.4
N	1,443	100

Note: The table shows descriptive statistics for families in our sample at the time of the first interview in 1998, when the designated child is 5 months old.

Table 2—: Parental behaviour

	Proportion of mothers who ...		
	Wave 1 5 months	Wave 2 17 months	Wave 3 29 months
Regularly checks on her child	71.7	47.8	31.9
Speaks spontaneously to her child	40.6	43.2	46.3
Answers to her child	45.0	46.8	57.4
Kisses and hugs her child	42.6	17.3	13.9
Screams toward her child	< 0.5	4.9	6.8
Is annoyed by her child	1.6	7.1	10.5
Reprimands her child	< 0.5	4.3	5.5
Supports her child progress	38.0	26.1	25.5
Organises play time	58.5	53.6	43.6
Gives pedagogical toys	68.2	59.0	43.7
Observations	1,443	1,443	1,443

Note: The table describes the behaviour of the respondents and their interactions with their children during the annual QLSCD interview. Behaviours are evaluated by the enumerator during the interview. Statistics are presented for the three first waves, when the designated child is 5, 17 and 29 months old.

3. Discovering latent parenting types

In the next step, the different features of parental behavior are summarized into interpretable behavioral types using a machine learning algorithm based on the latent Dirichlet allocation. This methodology developed by Blei, Ng and Jordan (2003) is a clustering algorithm for discrete data, which traditionally was meant to reduce the high dimensionality of text into an arbitrary number of topics specified by the user. Each parental action can be featured with difference importance in each type, and each parent can be a mixture of types.

The algorithm learns from the co-occurrence of counts through Bayesian learning. The idea is that if certain variables tend to appear together, they are likely to be linked to each other. In Appendix A we explain the technical details. For the sake of simplicity and interpretation, we settle on two types of parents. The final output of the algorithm is the distribution of actions for each type and the type distributions for each parent. With this information at hand, we can then relate parental types summarized into just two types to human capital accumulation.

3.1. Parenting types

We pool the three waves together and estimate the classification for that sample.[4] In Table 3 we display the absolute and relative occurrence of actions by the two types. The action that distinguishes the two types most in relative terms are supportive comments made by the parent to the child about its progress. While nearly two-thirds of parents of the active type make supportive comments about the progress of the child, this is the case for only 0.1% of parents of the laissez-faire type, i.e. they are 626 times more likely to do so. Similarly large differences exist for speaking to the child directly, which is done by 91% of the active parents compared to only 0.2% of the parents of the laissez-faire type.

The actions that are relatively more likely by laissez-faire parents are displaying annoyance, which is done by 3.5% compared to 3.3% of the active parents, and reprimanding the child,

[4]We could estimate a different classification for each wave separately as some actions might be more pertinent for different ages of the child, as is indicated by the distribution of actions in Table 2. However, the parental classification would not be comparable over time, which would pose other challenges for the rest of our analysis.

Table 3—: Classification of parental types

	Probability of occurrence		
	Type 1 "active"	Type 2 "laissez-faire'	Ratios
Supports her child progress	0.626	0.001	626 - 0.00
Speaks directly to her baby	0.908	0.002	454 - 0.00
Answers to her child	1.039	0.004	260 - 0.00
Kisses and hugs her child	0.513	0.002	256 - 0.00
Organises play time	1.011	0.070	14.4 - 0.07
Gives pedagogical toys	1.086	0.099	11.0 - 0.09
Regularly checks on her child	0.812	0.224	3.63 - 0.28
Scream toward her child	0.053	0.031	1.71 - 0.58
Is annoyed by her child	0.056	0.072	0.78 - 1.29
Reprimands her child	0.033	0.035	0.94 - 1.06

Note: The table describes the occurrence of behaviours for the two types found by the LDA algorithm. Behaviours are classified from what is the most different between the two types to what is the less. The last column displays the ratio of the probability for type 1 over the probability for type 2 (first number), and the ratio of the probability for type 2 over the probability for type 1 (second number).

which is done by 7.2% compared to 5.6% of the active parents. The distribution of actions across types suggests that what distinguishes parents is the richness of action by one type versus the lack of action by the other, hence the labels active and laissez-faire parents.

The LDA algorithm assigns to each parent a probability of being of type 1, the active type (and with the remaining probability they are of type 2, the laissez-faire type). The top panel of Figure 1 shows the distribution of the active type probability for the full sample and the bottom panel for each wave separately. We see a concentration two masses: one with a low probability of being of the active type (i.e. with a high probability of being of the laissez-faire type) and the opposite. Over time, parents tend to move from the active type to the laissez-faire type.

Figure 1. : Distribution of types

a) Pooled across all waves

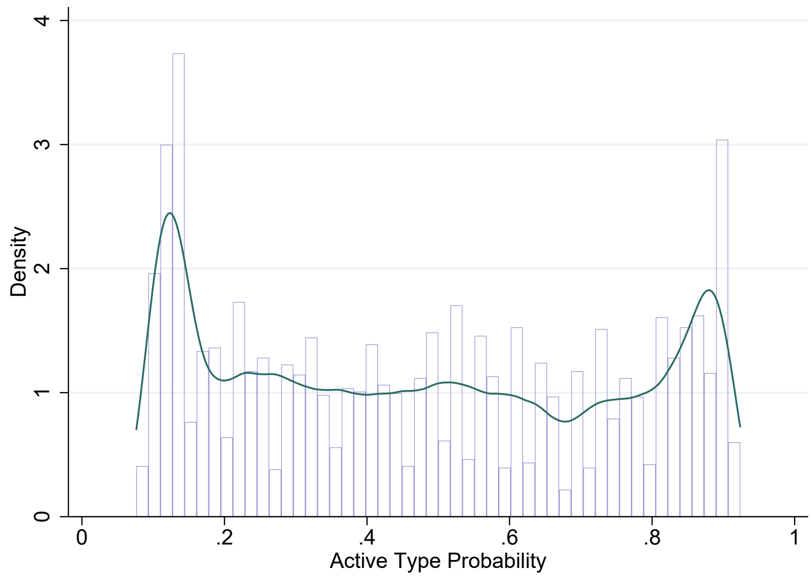

b) For each wave separately

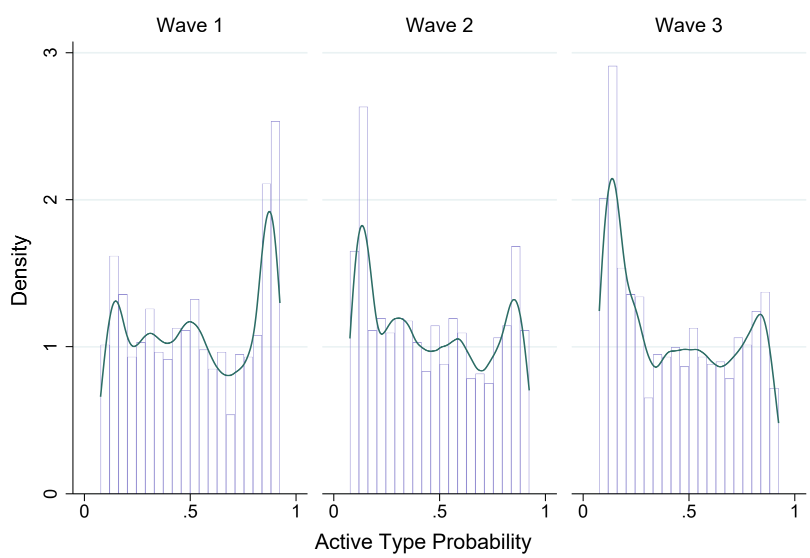

Notes: The transparent bars represent the binned probabilities of the probability of being an active type, while the solid line is the kernel density.

3.2. *Correlates and persistence of parenting types*

In Figure 2 we show the distribution of active types by maternal education. In the left panel we see that mothers with high school or less tend to be of the laissez-faire type with an average share of active types of 41.8%. In the middle panel we see that for mothers with some college education the distribution appears closer to bi-modal with an average probability of active types of 48.3%. Finally, in the right panel we see that amongst more educated mothers with a college degree, the average likelihood of being of the active type increases to 53.9%.

Figure 2. : Distribution of types by maternal education

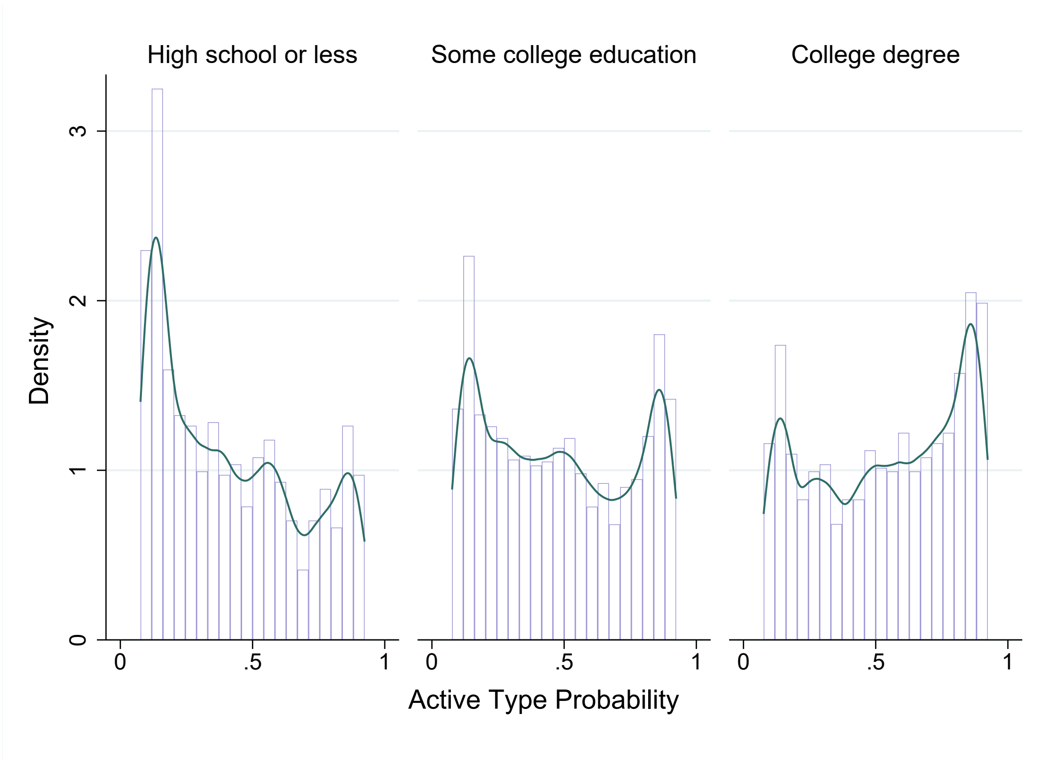

Notes: The transparent bars represent the binned probabilities of the probability of being an active type, while the solid line is the kernel density.

While the previous figure suggests that the likelihood of being an active mother is increasing in education, we take a more systematic look at the relationship between type and individual characteristics by regressing the probability of being an active type on age, education, poverty level, whether the parent is an immigrant, marital status, employment status, number of siblings, and the gender of child. In the first column of Table 4 we see the results for the pooled sample and in the following three columns for each age of the child, separately.

We see that parents with more than one child tend to be less likely to be of the active type. The probability of being active appears to be increasing in maternal age and education. While some of the coefficients vary, in general the direction of coefficients is very similar across waves. Maternal types reveal a considerable persistence as suggested by the correlations across waves exhibited in Table 5. Between wave 1 and wave 2 the correlation in types is 0.26, and between wave 2 and wave 3 it is 0.36. In fact, regressing individual fixed effects on parenting types achieves an R^2 of 0.52. We further breakdown the persistence in Table 6 in which we show the transition matrix between active types (defined as being of the active type with a probability above 0.67), an intermediate type (active type with a probability between 0.33 and 0.67), and the laissez-faire type (active type with a probability of less than 0.33). According to this matrix 38% (51%) of active (laissez-faire) mothers in wave 1 are of the same type in wave 2, and 42% (58%) of active (laissez-faire) mothers in wave 2 are of the same type in wave 3.

Table 4—: Active type probability and parental characteristics

	Probability of being of the active type			
	Pooled	Wave 1	Wave 2	Wave 3
Number of siblings (reference: no sibling)				
One sibling	-0.043***	-0.037**	-0.057***	-0.035**
	(0.011)	(0.016)	(0.016)	(0.016)
Two or more siblings	-0.060***	-0.039*	-0.105***	-0.035
	(0.017)	(0.023)	(0.023)	(0.023)
Household type (reference: two-parents family)				
Blended family	0.013	-0.005	0.035	0.010
	(0.018)	(0.025)	(0.024)	(0.023)
Single-parent household	-0.146	-0.160	-0.102	-0.176*
	(0.127)	(0.193)	(0.178)	(0.103)
Mother's age (reference: less than 25)				
25-29	0.024	0.053***	0.013	0.006
	(0.015)	(0.021)	(0.021)	(0.021)
30-34	0.055***	0.089***	0.046**	0.029
	(0.016)	(0.021)	(0.023)	(0.021)
35 and more	0.082***	0.100***	0.071***	0.075***
	(0.019)	(0.027)	(0.027)	(0.026)
Mother born outside Canada (reference: no)				
Yes	-0.035*	-0.061**	-0.042	-0.003
	(0.020)	(0.029)	(0.028)	(0.029)
Language spoken at home (reference: French)				
Other	-0.034**	0.034	-0.060***	-0.075***
	(0.014)	(0.021)	(0.020)	(0.021)
Mother education (reference: high school degree or less)				
Some college education	0.050***	0.042**	0.037**	0.073***
	(0.013)	(0.018)	(0.018)	(0.018)
College degree	0.089***	0.084***	0.079***	0.103***
	(0.015)	(0.022)	(0.022)	(0.022)
Parental working status (reference: two-parents: both work)				
Two-parents: one works	-0.001	0.014	-0.005	-0.012
	(0.013)	(0.019)	(0.019)	(0.019)
Two-parents: none work	-0.033	-0.004	-0.088*	-0.006
	(0.034)	(0.045)	(0.047)	(0.045)
Single-parent: works	0.136	0.127	0.091	0.189*
	(0.133)	(0.199)	(0.184)	(0.113)
Single-parent: does not work	0.151	0.147	0.102	0.204*
	(0.132)	(0.198)	(0.183)	(0.113)
Below poverty threshold (reference: no)				
Yes	-0.025	-0.055**	0.003	-0.022
	(0.016)	(0.023)	(0.022)	(0.022)
Constant	0.436***	0.455***	0.455***	0.400***
	(0.0159)	(0.022)	(0.022)	(0.022)
Observations	4,329	1,443	1,443	1,443
R-squared	0.050	0.070	0.060	0.056

Note: Each column presents the estimates of an OLS regression of active type probability on parental characteristics. The categories for missing values are also included in the regression but not shown in the table as they only concern a few individuals and are thus hard to interpret. Robust standard errors clustered at the family level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 5—: Correlation matrix of active-type probability across waves

	Wave 1	Wave 2	Wave 3
Wave 1	1.00	.	.
Wave 2	0.26	1.00	.
Wave 3	0.24	0.36	1.00

Note: The table displays the correlation between the active type probability variable in wave 1 and the one in wave 2, the active type probability variable in wave 2 and the one in wave 3, and the active type probability variable in wave 2 and the one in wave 3.

Table 6—: Transition matrix between binned types

(a) Between waves 1 and 2

	Wave 2		
Wave 1	Active	Intermediate	Laissez-faire
Active	0.38	0.37	0.25
Intermediate	0.27	0.36	0.37
Laissez-faire	0.17	0.32	0.51

(b) Between waves 2 and 3

	Wave 3		
Wave 2	Active	Intermediate	Laissez-faire
Active	0.42	0.36	0.23
Intermediate	0.27	0.35	0.37
Laissez-faire	0.14	0.28	0.58

Note: The first table presents the transition matrix between active types (defined as being of the active type with a probability above 0.67), an intermediate type (active type with a probability between 0.33 and 0.67), and the laissez-faire type (active type with a probability of less than 0.33) between wave 1 and wave 2. The second table presents the same transition matrix between wave 2 and wave 3.

4. Relating parenting types to children's outcomes

To test the relationship between parental type and the accumulation of children's cognitive skills, we use the results from an Imitation Sorting Task (IST) test conducted during each wave.[5] Here the sample size reduces to 1,121 children who took the IST test at 5, 17 and 29 months. Excluded children were sleeping or sick at the time the test was supposed to take place or the test was not fully completed. The test score in each wave is standardized with a mean of 0 and a standard deviation of 1.

In the first column of Table 7 we show the results of the pooled sample in which we regress the IST test score at each of the three stages on the probability of being an active parent and a constant. We find that moving from a laissez-faire to an active parent is associated with an increase in the IST score of 0.223 standard deviations. In the second column we add controls for parental characteristics and still find a highly significant positive association between the probability of being an active parent and test scores of 0.167 standard deviations. In the third column we control for parental fixed effects, thereby removing any constant heterogeneity across parents and children. Using this specification we find a strengthened association between being an active type and cognitive development with a highly significant coefficient of 0.338.

[5]The task comprises different situations in which the infant must grasp objects placed in front of him/her and place them in given containers. The task used in the ELDEQ is a variation of the Imitation Sorting Task developed by Uzgiris and Hunt (1975).

Table 7—: Active type probability and cognitive development

	Standardized IST score		
	(1)	(2)	(3)
Active type probability	0.223***	0.167**	0.338***
	(0.069)	(0.069)	(0.108)
Observations	3,363	3,363	3,363
R-squared	0.004	0.020	0.387
Family characteristics	NO	YES	NO
Family FE	NO	NO	YES

Note: Each column presents the estimates of an OLS regression of the child standardized IST score on her mother active type probability. Family characteristics (column 2) include family composition (number of siblings, household type), maternal characteristics (age, whether born outside Canada, educational attainment), parental working status, language spoken at home, and whether family is below poverty threshold. They are described in more details in Table 1. Robust standard errors clustered at the family level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

5. Conclusion

Human capital accumulation is one of the most important fundamentals of productivity and innovation. However, estimating human capital production functions is riddled with complications including the high dimensionality and potentially non-linear relationships between parental investments. In this paper we provide a new way to summarize parental investments adopted from computational linguistics. We use an unsupervised machine learning model, the latent Dirichlet allocation, to classify parents into two types. The resulting types can be interpreted as active parents who encourage their children and express their affection, versus laissez-faire parents who do not interact much with their children.

We show that these two types relate systematically to parental characteristics, i.e. mothers with higher education tend to be morel likely to be of the active type. Moreover, we show that children of more active parents tend to achieve higher levels of human accumulation. While we cannot establish a causal relationship between parenting types and outcomes due to the nature of the data, we are optimistic that future studies including natural experiments or randomized control trials can make use of the proposed methodology to classify parents into types based on their their actions. Another advantage of the approach is that this can be done with an extremely large set of actions or even detailed time use data.

REFERENCES

Agostinelli, Francesco, Matthias Doepke, Giuseppe Sorrenti, and Fabrizio Zilibotti. 2020. "It takes a village: the economics of parenting with neighborhood and peer effects." National Bureau of Economic Research.

Attanasio, Orazio, Costas Meghir, and Emily Nix. 2020. "Human capital development and parental investment in india." *Review of Economic Studies.*

Bandiera, Oriana, Stephen Hansen, Andrea Prat, and Raffaella Sadun. 2017. "Ceo behavior and firm performance." National Bureau of Economic Research.

Blei, David M, Andrew Y Ng, and Michael I Jordan. 2003. "Latent Dirichlet Allocation." *Journal of Machine Learning Research*, 3(Jan): 993–1022.

Cobb-Clark, Deborah A, Nicolas Salamanca, and Anna Zhu. 2019. "Parenting style as an investment in human development." *Journal of Population Economics*, 32(4): 1315–1352.

Cunha, Flavio. 2015. "Subjective rationality, parenting styles, and investments in children." In *Families in an Era of Increasing Inequality*. 83–94. Springer.

Cunha, Flávio, James J Heckman, and Susanne M Schennach. 2010. "Estimating the technology of cognitive and noncognitive skill formation." *Econometrica*, 78(3): 883–931.

Del Boca, Daniela, Christopher Flinn, and Matthew Wiswall. 2014. "Household Choice and Child Development." *The Review of Economic Studies*, 81(1): 137–185.

Del Boca, Daniela, Christopher J Flinn, Ewout Verriest, and Matthew J Wiswall. 2019. "Actors in the Child Development Process." National Bureau of Economic Research.

Doepke, Matthias, and Fabrizio Zilibotti. 2017. "Parenting with style: Altruism and paternalism in intergenerational preference transmission." *Econometrica*, 85(5): 1331–1371.

Doepke, Matthias, and Fabrizio Zilibotti. 2019. *Love, money, and parenting: How economics explains the way we raise our kids.* Princeton University Press.

Doepke, Matthias, Giuseppe Sorrenti, and Fabrizio Zilibotti. 2019. "The economics of parenting." *Annual Review of Economics*, 11.

Draca, Mirko, and Carlo Schwarz. 2018. "How polarized are citizens? Measuring ideology from the ground-up." *Measuring Ideology from the Ground-Up (April 2, 2018).*

Hoffman, Matthew, Francis R Bach, and David M Blei. 2010. "Online learning for latent dirichlet allocation." 856–864.

Pedregosa, Fabian, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. "Scikit-learn: Machine learning in Python." *the Journal of machine Learning research*, 12: 2825–2830.

Uzgiris, Ina C, and Joseph Hunt. 1975. "Assessment in infancy: Ordinal scales of psychological development."

Appendix A: Latent Dirichlet allocation

Adapting the technical terms from Blei, Ng and Jordan (2003) for text and applying to our objective, the corpus of behavioral actions D is composed of parents w of actions. A behavioral type is a probability distribution over all actions. The assumed underlying process with which types generate actions is by drawing θ from a Dirichlet distribution with hyperparameter α. Then for each action n of all actions N, one chooses a type from z_n. After that an action w_n is chosen for the corresponding type z_n from a Dirichlet distribution with hyperparameter β.

Written formally, the generative process of actions is expressed as the following joint distribution

$$p(\beta, \theta, z, w_d) = \prod_{i=1}^{k} p(\beta_i) \prod_{d=1}^{D} p(\theta_d) (\prod_{n=1}^{N} p(z_{d,n}|\theta_d) p(w_{d,n}|\beta, z_{d,n})).$$

Given the corpus of actions, the task of the algorithm is to infer the type-specific action distribution and the parent specific type distribution. So the posterior distribution of the latent variables is given by

$$p(\beta, \theta, z|w_d) = \frac{p(\beta, \theta, z, w_d)}{p(w_d)}.$$

In order to infer the marginal distribution $p(w_d)$, which can be done through approximation using Gibbs sampling, or Variational Kalman Filtering and Variational Wavelet Regression, we rely on the Stata implementation developed by Draca and Schwarz (2018). Draca and Schwarz (2018) use the inference algorithm developed by Hoffman, Bach and Blei (2010) and implemented by Pedregosa et al. (2011). As is the case in Draca and Schwarz (2018), the assumption of the independence of responses does not strictly hold in our approach. If an action has been recorded the same action is not recorded again for the same person. They discuss in detail why the inference of LDA is nonetheless still valid.