SCHOOL CLOSURES, PARENTAL LABOR SUPPLY, AND TIME USE

Enghin Atalay Ryan Kobler Ryan Michaels

November 2023

Abstract

The closure of schools to in-person instruction during the COVID-19 pandemic posed a unique shock to parents. This paper re-examines the effect of schooling mode on parental labor supply. The effects are undetectable using a full suite of controls for unobserved heterogeneity, which can be motivated by the failure of more parsimonious models to pass simple placebo tests. Even abstracting from such controls, though, a shift from fully virtual to in-person implies an increase in hours worked of 2 to 2.5 hours per week. We present a simple model of parental time allocation and child development to formalize why these estimates appear unexpectedly small. We then introduce telework and nonparental care into the theory, demonstrate that these features can support realistic labor supply outcomes, and illustrate how our estimates in turn discipline the inference of salient structural parameters. Evidence from time use diaries indicates that telework did support both market work and childcare, chiefly among parents with college degrees. Time use data and other surveys also provide suggestive evidence of the increased utilization of nonparental care.

JEL codes: J21, J22, J48

Keywords: Labor supply, childcare, pandemic, in-person learning

---

Atalay, Kobler, and Michaels: Research Department, Federal Reserve Bank of Philadelphia. The views expressed in this paper are solely of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

Beginning in March 2020, U.S. schools switched to remote instruction, and many did not reopen for consistent in-person instruction for a year. The closure of schools to in-person instruction posed a unique shock to parents. As Goldin (2022) notes, there were significant concerns at the time that the adoption of remote instruction would upend the careers of working parents. However, comparisons of labor force outcomes across adults with and without school-age children do not point to a dramatic change in parents' relative working time (Goldin, 2022; Furman et al., 2021).

These initial findings raise potentially interesting questions for theories of labor supply and time use more generally. How did parents ease the trade-off between market work and childcare? Put another way, on what other margins did parents adjust? And what might these decisions imply about the preferences, technologies, and constraints shaping time allocation decisions?

As a first step toward pursuing these questions, our paper begins by revisiting evidence on the effect of remote instruction on parental labor supply. Following leading work by Garcia and Cowan (2022) and Hansen et al. (2022), we link adults' working time to the local schooling mode. As detailed in Section 1, we measure changes in the prevalence of in-person instruction using Parolin and Lee's (2021) estimates of visits to school campuses derived from SafeGraph's mobile phone location data. The ratio of visits during some month in the pandemic era to visits during the same month in 2019 is interpreted as the in-person share of instruction. These estimates, aggregated to the county or higher geographic unit, can then be combined with observations on individual working time in the Current Population Survey (CPS).

We apply these data to estimate the response of individual working time to variation in the local in-person share. The potential endogeneity of school policy poses an immediate challenge. For example, one concern is that both school policy and parents' labor supply trend in the same direction as the general public's preference for a return to "normal" activity. Such preferences are not directly or fully observed.

As we discuss in Section 2, a common way of addressing this challenge is to leverage variation in working time across adults with and without children (see Garcia and Cowan, 2022; Hansen et al., 2022; Heggeness and Suri, 2021). This approach is grounded in the simple observation that school policy should have a direct effect only on parents of school-age children. Therefore, adults without children can potentially serve as a control group. The identifying assumption is that parents' relative hours worked—relative, that is, to hours worked of childfree adults—reflect only school policy, controlling for observables.

While this strategy effectively differences out market-wide factors, it potentially neglects systematic differences between parents and childless adults. Parents sort into different jobs and firms (Adda et al., 2017; Kleven et al., 2019); hold different political views and public policy preferences (Elder and Greene, 2013; Kerry et al., 2022); and exhibit different degrees of risk aversion in general and regarding COVID-19 specifically (Görlitz and Tamm, 2020; Elder and Greene, 2021). This heterogeneity may translate into different decisions with respect to economic activity as well as alternative perspectives on the path of school policy, which is accountable to parents. Therefore, we propose in Section 2 to augment the regression with parental status fixed effects intended to capture differences in parents' circumstances and preferences over time (but common across space) and across space (but fixed over time).

The results presented in Section 3 indicate that the estimated effect of school policy is indeed sensitive to the choice of controls for unobserved heterogeneity. According to the most parsimonious specification, a switch from virtual to in-person instruction implies an increase in hours worked of 0.5 per week. This result is, however, unstable across sub-samples; if we restrict attention to the 2020-21 school year, weekly hours rise by two. This instability is resolved by controlling for parental status-by-time effects, revealing a stable coefficient of around two hours per week. Thus, controlling for unobserved heterogeneity in this dimension recovers a larger response. Controlling for spatial heterogeneity, though, has the opposite effect, and dramatically so. After including parental status-by-area effects, the association between in-person shares and hours worked is eliminated.

To assess these competing specifications, Section 3 considers a simple placebo test. This exercise examines, specifically, if differences in pandemic-era school policies across areas may reflect longer-term sources of parental heterogeneity (across space). If so, average pandemic-era in-person shares should predict parents' relative hours worked in the *pre*-pandemic period. We confirm this is so, and, moreover, the size of this correlation is on par with the size of the estimated effects of school policy on pandemic-era data.

While we see this evidence of a null effect as compelling, estimates derived from more parsimonious models are still instructive. As we discuss in Section 1, the predominant source of variation in school policy seems to be spatial. As a result, a null effect does not necessarily rule out an allocative effect of policy but instead indicates there is insufficient variation to tease it out. At the same time, the placebo results suggest one is likely to overstate the true impact of school policy by neglecting this form of parental heterogeneity. Accordingly, when abstracting from spatial heterogeneity, one can arguably interpret the estimates as upper bounds on the true effect.

In this spirit, we present a battery of results based on this simpler specification. A few findings stand out. First, labor supply responses are slightly higher among mothers than fathers, but one could not reject equality. The gender gradient is somewhat larger, though, among college graduates; the labor supply of fathers with a college degree is essentially inelastic with respect to in-person share. Second, parents of younger school-age children (those with children aged 5-9) seem to be the most responsive to school format changes. Third, labor supply responses vary little by marital status but do vary notably within the unmarried. Estimates for the latter aggregate two different results: labor supply is relatively responsive among (male and female) lone-adult parents—work hours increase by 4.5 per week when in-person instruction is reinstated—but not among the unmarried in co-residential arrangements with other adults.

Even when this analysis does uncover significant effects, though, they seem rather modest. Consider again a parent whose school district transitions away from a virtual format. This change reintroduces over 30 hours of in-person instruction per week, and yet the estimated increase in labor supply is a small fraction of this, even among groups (such as lone adults) who are arguably most exposed to the policy. The remainder of the paper explores what margins of adjustment may help account for these results and situates them within a simple model of parental time use.

To this end, we next examine broader time use patterns. drawing on the American Time Use Survey (ATUS). We uncover several findings. First, we observe little adjustment along any major time use margin—neither leisure, market work, nor home production—to variation in in-person instruction shares. Second, telework was likely one means by which some parents insulated their schedules from pandemic disruptions. Our estimates suggest that a shift from in-person to virtual school formats led college-educated parents to spend 6 more hours per week working from home while simultaneously looking after their children. We observe no telework response among the noncollege educated, consistent with the observed divide in telework opportunities by education (Mongey et al., 2021). Third, nonparental care was used more intensively after the suspension of in-person instruction. We find that, in this event, over-60 respondents—a group likely to include many grandparents—allocated over 3 more hours per week to the care of *others'* children.[1] Because of the small sample size of the ATUS, these estimates are subject to considerable uncertainty. Still, the results point to two potentially promising explanations of the labor supply findings.

In Section 5, we view these results through the lens of simple models of parental time allocation. The simplest setup deliberately neglects telework and nonparental care to formalize the sense in which the estimated labor supply effects are, at first glance, surprisingly small. Following Berlinski et al. (2023), a parent in the model values consumption, leisure time, and child development. To start, we assume a child's development is a function of two arguments: the parent's supervision and a form of publicly provided supervision, e.g., in-person class time. In addition, a child must always be supervised by a parent or by school. In this context, a decline in publicly provided supervision leads the parent to substitute time toward childcare (and away from leisure and market work). We show that the model predicts counterfactually large labor supply

[1] Unfortunately, the ATUS does not report the identity, or even the educational attainment, of the parent of the child who received care from the over-age-60 respondent.

effects. In fact, if the change in school policy is seen as temporary and if the parent can smooth the family's consumption in the meantime, then labor supply will drop one for one with in-classroom time. Even in the polar case where the family lives "hand to mouth", the calibrated model predicts a labor supply response many times larger than we observe.

We then amend this baseline setup to illustrate the potential roles for telework and nonparental care. First, we introduce a novel "multi-tasking" technology to capture the idea that teleworking enables parents to carry out, to an extent, multiple tasks at the same time, e.g., working while simultaneously supervising children. The technology is indexed by just a single (new) parameter, and we derive the mapping from the latter to the labor supply response. Second, noting that many parents did not have access to a telework opportunity, we next consider a margin of adjustment omitted from the baseline model, namely, nonparental care. We show that our labor supply findings are consistent with parental and nonparental care being strong substitutes in child development (Berlinski et al., 2023). This section concludes by highlighting the broader implications of this substitutability for public policy and cyclical hours dynamics.

Related research. Our paper contributes to several strands of research. First, our analysis of CPS data extends earlier efforts by Garcia and Cowan (2022) and Hansen et al. (2022), who study the link between local in-person instruction shares and individual hours worked. Our placebo results illustrate why often-used specifications, which neglect some sources of heterogeneity, are likely to yield upper bounds on the labor supply response. Consistent with this observation, estimates from such specifications are generally near the top end of the range of causal estimates of childcare availability on parental labor supply. For instance, in an analysis of the introduction of public kindergarten, Gelbach (2002) and Cascio (2009) find similar or slightly smaller estimates for unmarried mothers but notably smaller (or null) responses for married mothers.[2] Studies of contexts outside the U.S. yield more varied results. Longer mandated school hours yield relatively large gains in maternal hours in Padilla-Romo and Cabrera-Hernández (2019) and Berthelon et al. (2023), but more modest increases in Contreras and Sepúlveda (2017). Increased after-school care is found to have no net employment impact in Felfe et al. (2016).[3]

Next, our analysis of the ATUS contributes to a growing research agenda on telework. Pabilonia and Vernon (2023) document that take-up of remote work increased at the onset of the pandemic, especially for mothers of children under the age of 13. When working from home, parents spent a large share of their time on secondary childcare activities. Atalay (2023) shows that these shifts were more pronounced for parents with a college degree (see also Cowan, 2023).

[2] Cascio argues that the availability of other forms of nonparental care mitigated the impact of public kindergarten.
[3] A related strand of research documented changes in hours worked in the months immediately after the onset of the pandemic. Some of this research found substantial movements in parental hours (Alon et al., 2020; Heggeness, 2020), whereas others found more muted responses (Lozano-Rojas et al. 2020; Barkowski et al., 2021). Our analysis will span all of 2020-21 and with more of a focus on the period beginning with the fall 2020 school year.

Our results broadly echo these earlier findings on the incidence of remote work and caregiving during the pandemic. We extend this research by more precisely linking parental time use patterns to local in-person shares.

Finally, we connect pandemic-era research on school closures to economic models of parental time use. We show analytically how results in our paper and elsewhere can inform theories of parental investments and adolescent development and illustrate their broader implications for policy interventions and labor market dynamics. In addition, we offer a new means to formalize telework and, thus, extend this branch of theory to incorporate the widespread use of remote work. We view our analysis, which draws out lessons from the data within simple models, as complementary to the structural estimation of much richer models (see Del Boca et al., 2014, and Berlinski et al., 2023).

1. Data

This section reviews the data on school policies, labor supply, and other controls used in the subsequent regression analysis.

1.1 School policies

The pandemic prompted almost all school districts to shift toward remote instruction in March 2020. Although many retained this format to start the 2020-21 school year, modes of instruction did begin to diverge then—even across neighboring counties. For instance, the Atlanta district in Fulton County operated strictly remotely, whereas Forsyth County, just 40 miles north, made in-person instruction available to all students (Education Week, 2020).

The variation in school reopening plans spurred the creation of numerous schooling mode trackers, which aim to document the predominant mode of instruction in school districts. A few prominent sources include the American Enterprise Institute's (AEI) Return2Learn database, Burbio's School Reopening Tracker, and the COVID-19 School Data Hub (CSDH). These trackers vary with respect to the breadth of their coverage (e.g., the number of school districts in the sample); level of detail (i.e., grade-level v. district-wide outcomes); and data collection methods (i.e., web scraping v. school- and district-level surveys). The in-person instruction shares can vary widely across the trackers, which suggests that the different choices of methodology and sampling can meaningfully shape the results (Kurmann and Lalé, 2023).

Alternatively, some recent research has adopted a more indirect, but also more easily quantifiable, proxy of on-site instruction, namely, the volume of "foot traffic" on school campuses (Garcia and Cowan, 2022; Hansen et al., 2022). The source of the underlying data is SafeGraph, which obtains GPS data from individual mobile phones by pinging certain apps. The location data enable SafeGraph to track the number of visits to over 7 million points of interest (POI) in the U.S.

We will draw specifically on Parolin and Lee's (2021) tabulations of SafeGraph data. For each POI identified as a public school, Parolin and Lee calculate the percent change in visits between year $y \geq 2020$ and month m relative to the same month m in 2019.[4]

Parolin and Lee present two county-level measures. One is a straightforward average of the percent change in visits across schools (in each month). The other is derived by first categorizing a school as "closed" in some month m (and year $y \geq 2020$) if the number of visits to that school at that time is down by at least 50 percent relative to month m in 2019. Each school (in each month) is assigned a one if it is categorized as closed and zero otherwise, and Parolin and Lee report the average over this binary indicator. The complement of this measure—that is, one minus the Parolin and Lee figure—can be interpreted, roughly, as the in-person instruction share. The latter has been favored in the related literature and will be our default measure of school policy.

We see several advantages in the SafeGraph data. First, it is arguably the most comprehensive source of data in this literature, covering over 100,000 schools and virtually every county during the 2020-21 and 2021-22 school years. In addition, the use of mobile phone data naturally accommodates heterogeneity in learning modes. Within a district, some schools—and, within those schools, some students—may attend on-site while others operate predominantly remotely. Other schooling-mode trackers would classify the district according to one of a few coarse, discrete formats, such as "hybrid" or "virtual," whereas SafeGraph's data implicitly aggregates these modes into a single estimate of the change in on-site activity. In this sense, SafeGraph offers both a breadth of coverage and a level of precision that is unique.

Still, mobile phone data are not immune to measurement error. The number of mobile phone pings may not necessarily be proportional to the number of students engaged in on-site instruction. Suppose, for instance, that faculty at a primary school are asked to, or prefer to, work in their classrooms when they teach virtual lessons (see Cohen, 2020; and Jung, 2020). This policy attenuates the decline in foot traffic even if on-site instruction is prohibited. Clearly, instruction-mode trackers based on published school district policy would not commit this error.

Nevertheless, our empirical analysis treats SafeGraph—and, specifically, Parolin and Lee's figures—as our baseline. However, later in Section 3, we contrast SafeGraph-based results with estimates derived from two instruction-mode trackers, namely, Burbio and CSDH.

Geographic variation in in-person shares. Although Parolin and Lee's school-level estimates cover the more than 3,000 U.S. counties, our other data sources do not offer this same breadth and detail. Crucially, the Current Population Survey, our source on parents' hours worked, neither discloses school districts nor universally reports the respondent's county. Over the two

[4] Parolin and Lee exclude private schools because, for their analysis, they link their school-level estimates to student demographic data available only for public schools.

(calendar) years 2020-21, the CPS identifies only 280 counties. Although the latter are relatively large, one's county is not disclosed for 60 percent of (adult) survey respondents. Fortunately, though, the CPS identifies the metropolitan statistical area (MSA) for almost 60 percent of those with no reported county. A respondent's state is always provided.

In view of these constraints, we apply a three-step method to aggregate SafeGraph data and integrate it into the CPS (see Hansen et al., 2022). First, if a county is reported in the CPS, we assign its respondents the county in-person share from Parolin and Lee. Second, if a collection of counties is not identified in the CPS but does belong to a disclosed MSA, we aggregate Parolin and Lee's estimates across these counties and assign the mean to CPS respondents in that MSA for which no county is reported. Finally, we aggregate Parolin and Lee's estimates among counties within a state that are not reported in the CPS and do not belong to a reported MSA. The mean among these counties is assigned to CPS respondents in the state for which no county or MSA identifier is provided. In total, by aggregating within MSA where feasible and within state where necessary, we identify 198 more areas to reach a total of 478.[5] This strategy maximizes the use of the Parolin and Lee data.

Figure 1 illustrates the variation in school policies implied by the SafeGraph data. For each of the 478 areas, the figure plots the average in-person instruction share in September-December 2020 along the x-axis and the average share in January-May 2021 along the y-axis. There are a few points to note. First, there are significant differences across areas. In each of the two semesters, in-person shares span a wide range from 0.2 to 1. These regional differences are, to some extent, persistent: almost half of the areas lie within 10 percentage points of the 45-degree line, that is, the in-person share shifted by less than |0.1| across semesters. Among the other half of the areas, though, many observations lie well off the 45-degree line, which indicates substantial variation in school policy *within* region. The latter variation generally reflected differences in the timing of reinstating in-person instruction in spring 2021; few regions cut in-person learning in 2021.

What might account for the differences in school policies illustrated in Figure 1? And are any of these sources of variation likely to shape labor supply? Clearly, one possible source is the spread of COVID-19: if the threat of infection and fatality were to recede, we may see increases in in-person instruction and labor supply, even if the former has no causal effect on the latter.

[5] The additional local areas include 151 MSAs, or subsets of MSAs. If a county is reported in the CPS, it is not included in our construction of an MSA-based local area. The remainder of local areas comprises data from 47 states where we observe CPS respondents who do not belong to a disclosed county or MSA. The reason this step captures data from only 47 states, rather than 51, is that, for a handful of very small states, all survey respondents live in a disclosed county or MSA.

Figure 1: In-Person Shares in 2020-21 School Year

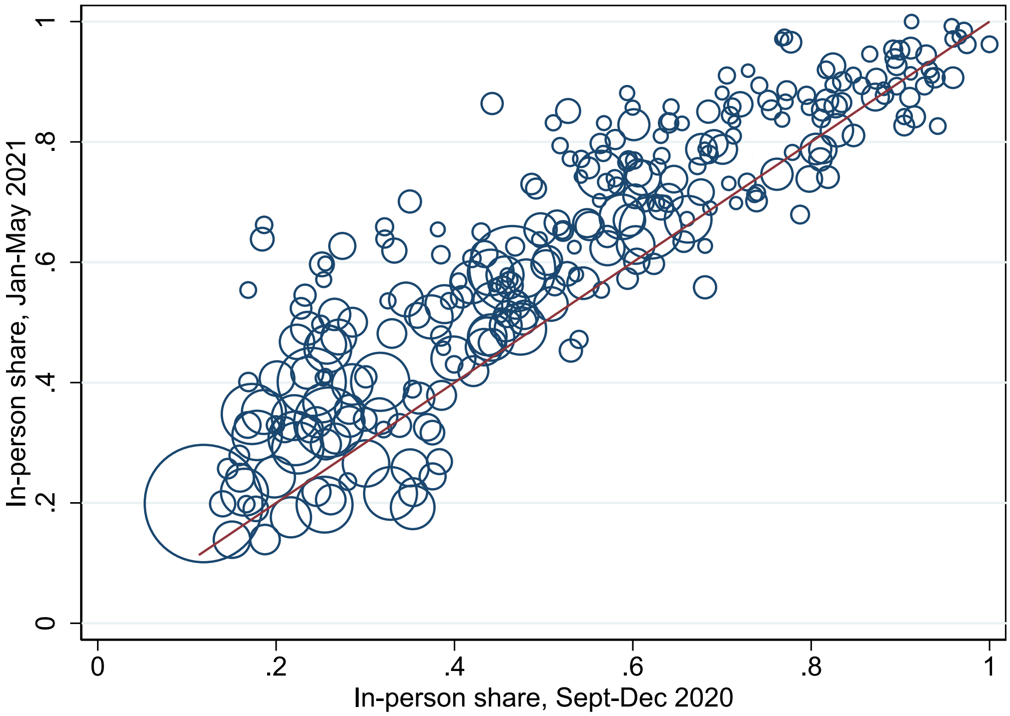

Note: This figure plots the average in-person share in September to December 2020 (x-axis) compared to January to May 2021 (y-axis). The size of each circle is proportional to the total enrollment in the geographic area.

In fact, the link between school policy and COVID-19 case counts is surprisingly modest. Online Appendix A fleshes out the evidence on this point, though it is consistent with what many related papers report (see below). Importantly, this result applies to variation (in school policy and COVID-19 cases) both *within* and *across* regions. We suspect that, within a region, monthly changes in case counts are only weakly correlated with changes in policy because the latter had to be set well in advance of implementation. For example, Prince George's County (Maryland) announced in mid-July 2020 that it would not consider a return to in-person instruction before February 2021. Just to its south, Fairfax County (Virginia) announced that, while it anticipated a shorter period of virtual instruction, it would not reinstate on-site instruction until at least the start of November. (In each county, COVID-19 cases had been on the decline throughout the summer.) These examples suggest that current school policy was partially predetermined and, therefore, likely to be somewhat insensitive to changes in the state of the pandemic.

Instead, school policy appears to be shaped by regional political forces. Partisan affiliation and, more concretely, the degree of support for Donald Trump were significant predictors of school policy. The strength of teacher unions also helps account for variation in in-person shares.[6] These factors would seem to reflect long-held local preferences and norms, which in turn may be correlated with labor market activity independent of school policy. If so, regional variation in policy may proxy for other employment-relevant factors. We return to this point in Section 2.

1.2 Summary of sample

We draw on several data sources for our main regressions (in addition to the aforementioned measures of on-site instruction). Labor supply and worker demographics are taken from the monthly Current Population Survey (Flood et al., 2022). We typically measure labor supply as weekly hours of work in the survey reference week but also report results where the outcome is employment status (in the reference week). Other variables measure the state of the pandemic and public health policy responses. We draw on county-level data on COVID-19 cases and deaths published weekly by Johns Hopkins Coronavirus Resource Center (Dong et al., 2020).[7] These data are aggregated up to the monthly frequency and to the local geographic areas described above. We also use Kaiser Family Foundation measures of government mitigation policies, such as capacity limits on restaurants and bars.[8]

Table 1 reports means for many of the variables that will be used in our regressions. The averages are presented for several different subgroups of the population, distinguished by sex, age, and location. We report results here (and later in the paper) separately for men and women. Table 1 considers two age groups, 21 and over, and the narrower range of 21-59. In addition, the table reports results for parents of school-age children. (The ages of parents are unrestricted, but nearly all fall within the range 21-59.) Finally, tabulations are shown for CPS-reported counties as well as the full sample of local areas described above. As discussed later, our regression sample consists of all areas but restricts attention to ages 21-59. It is instructive, though, to contrast our preferred sample to the alternative groups in Table 1.

A few patterns in the data are particularly noteworthy, if not necessarily unexpected. First, consider the differences across age ranges and parental status for a given gender and geographic coverage. The sample of adults of all ages is (naturally) older, has fewer kids in the home, is less racially and ethnically diverse, and works less than the other two. In other words, this subsample

---

[6] For results on partisanship and union strength, see Grossman et al. (2021), Hartney and Finger (2021), and Marianno et al. (2022). Online Appendix A reports that the interaction between the latter and COVID-19 cases are statistically significant predictors of instruction format but still account for a very limited share of the variance in in-person shares.
[7] These data can be found at `https://github.com/CSSEGISandData/COVID-19/tree/master/csse_covid_19_data/csse_covid_19_time_series`. Accessed August 2, 2023.
[8] These data can be found at `https://github.com/KFFData/COVID-19-Data/tree/kff_master/State%20Policy%20Actions`.

is observationally quite different than the "treated" group, namely, the parents of school-age children. By contrast, the full sample of adults ages 21-59 is very similar to parents along nearly

Table 1: Summary Statistics

Variable	CPS Counties			All Local Areas		
Women						
	Age ≥ 21	21 – 59	Parents	Age ≥ 21	21 – 59	Parents
Weekly hours	19.221	24.791	23.745	19.255	24.978	24.185
Employment	0.519	0.662	0.645	0.521	0.666	0.655
Age	49.941	39.793	41.122	50.074	39.810	40.650
Kids in home	0.219	0.319	1.000	0.225	0.330	1.000
Bachelor or more	0.405	0.442	0.432	0.376	0.412	0.407
White	0.739	0.717	0.715	0.768	0.743	0.744
Black	0.141	0.152	0.151	0.134	0.145	0.141
Hispanic	0.200	0.233	0.278	0.160	0.192	0.230
Foreign born	0.246	0.255	0.315	0.188	0.203	0.254
Married	0.510	0.514	0.703	0.528	0.534	0.703
Resides in city center	0.342	0.358	0.318	0.286	0.304	0.270
Mo. cases / 100,000	691	686	694	711	706	710
In-person instruction	0.586	0.582	0.590	0.647	0.642	0.650
Number of obs.	314,807	201,996	66,131	763,440	482,202	165,914
Men						
	CPS Counties			All Local Areas		
	Age ≥ 21	21 – 59	Parents	Age ≥ 21	21 – 59	Parents
Weekly hours	25.951	31.658	35.559	26.106	32.124	36.225
Employment	0.640	0.772	0.845	0.639	0.776	0.851
Age	48.474	39.429	43.800	48.748	39.588	43.363
Kids in home	0.195	0.265	1.000	0.200	0.274	1.000
Bachelor or more	0.386	0.385	0.420	0.351	0.350	0.390
White	0.757	0.735	0.749	0.785	0.764	0.779
Black	0.127	0.137	0.118	0.119	0.129	0.105
Hispanic	0.209	0.242	0.279	0.169	0.201	0.235
Foreign born	0.244	0.255	0.338	0.187	0.204	0.274
Married	0.556	0.500	0.854	0.569	0.516	0.850
Resides in city center	0.345	0.362	0.303	0.286	0.307	0.254
Mo. cases / 100,000	689	685	690	711	707	708
In-person instruction	0.585	0.580	0.591	0.647	0.642	0.648
Number of obs.	285,048	191,348	53,742	701,227	462,290	136,596

Note: "CPS Counties" refers to the sample of counties that are recorded in the Current Population Survey. "Parents" are adults with at least one child between the ages of 5 and 17 in the household. Monthly cases / 100,000 refer to the contemporaneous number of COVID-19 cases in the local area of the respondent in the survey month. In-person instruction refers to the share of respondents in a local area that is open to in-person instruction in the survey month.

all dimensions (with marital status the obvious exception). Next, individuals sampled in CPS counties are relatively urban, educated, and ethnically diverse (for given age and gender). Notably, these counties also elected to operate schools in person less often. Thus, the use of all local areas captures a sample more broadly representative of parents and school policies. Finally, well-known differences in labor market participation and marriage rates between mothers and fathers are apparent in the table. The labor supply of single mothers will be an important topic in the analysis below.

2. Empirical framework

Our aim is to examine the effect of in-person instruction on parental labor supply. In this context, the potential endogeneity of instruction format is the most significant concern for estimation. In this section, we discuss a series of controls aimed at mitigating the endogeneity problem.

We first consider the empirical strategy adopted in much of the related literature (Garcia and Cowan, 2022; Heggeness and Suri, 2021; Collins et al., 2021). This approach leverages differences in hours worked across adults with and without children to identify the effect of school policy. The argument is that, even if policy is endogenous to the overall state of the labor market, it is arguably (as good as) random with respect to parents' *relative* labor market experiences.

This approach is formalized as follows. Denote the presence of one's own children in the home in month t by the indicator $\mathbb{k}_{it} = \{0,1\}$. The latter equals one if survey reference person i reports that he or she has children of school age in the residence. Next, let p_{at} denote the in-person instruction share in area a.[9] The effect of interest is, specifically, the parental labor supply response to variation in p_{at}. Accordingly, we adopt the estimating equation,

$$h_{iat} = \alpha\mathbb{k}_{it} + \beta p_{at} + \psi p_{at}\mathbb{k}_{it} + \boldsymbol{\gamma}'\mathbf{x}_{it} + \chi_a + \tau_t + \varepsilon_{iat}, \quad (1)$$

where h_{iat} is hours worked of individual i in area a in month t. The vector $\mathbf{x}_{it}$ captures additional individual-level controls to be described in the next section (and $\boldsymbol{\gamma}$ is a conformable vector); χ_a is an area fixed effect; and τ_t is a month fixed effect.[10] The key parameter in equation (1) is ψ, which measures the parental hours response to a unit difference in the in-person share.

Crucially, ψ can be estimated consistently even if school policy is endogenous to local area trends. The latter variation is "soaked up" by p_{at} and reflected in β, which measures the response

[9] In practice, this share varies *within* area, e.g., across school districts. Nevertheless, OLS yields consistent estimates so long as p_{at} correctly measures the mean of district-level shares.
[10] We have replaced τ_t with month-by-Census division effects, but this added granularity makes little difference.

that is common across *all* adults. A $\beta \neq 0$ is a natural outcome if, for instance, policy generally follows in line with a return to "normalcy", which shapes market-wide labor supply and demand.

The key assumption underlying equation (1) is that adults with and without children do not have systematically different labor supply preferences or face systematically different labor demands. From this assumption, it follows that $\varepsilon_{iat} \perp \mathbb{k}_{it}$. In a second specification, we partially relax this restriction. Consider the estimating equation,

$$h_{iat} = \beta p_{at} + \psi p_{at}\mathbb{k}_{it} + \zeta_a\mathbb{k}_{it} + \theta_t\mathbb{k}_{it} + \boldsymbol{\gamma}'\mathbf{x}_{iat} + \chi_a + \tau_t + \omega_{iat}, \tag{2}$$

which introduces two new fixed effects. The month effect, θ_t, captures *parent-specific* factors behind hours worked that are common across areas but vary over time, whereas the area fixed effect, ζ_a, captures mean level differences across space. Controlling for these effects, the identifying assumption is now that the labor supply incentives facing parents and childless adults do not differ systematically *over time within* a local area.

The controls in equation (2) serve to narrow the channel through which school policy acts on hours worked. Average regional variation in parents' relative hours worked is captured by $\zeta_a\mathbb{k}_{it}$, whereas temporal variation that is common across areas is picked up by $\theta_t\mathbb{k}_{iat}$. Thus, estimation of equation (2) recovers a significant effect of school policy only to the extent that parents' relative hours worked vary over time with the in-person share in their area.[11] By contrast, equation (1) draws on both the within- and across-area comovement of school policies and parents' relative hours. Thus, equation (2) offers potentially more credible identification of the effect of school policy but at the cost of statistical power.

3. Estimates from the CPS

In this section, we report estimates from the regression models just discussed. After we specify our sample and list of controls, we present our baseline estimates in Section 3.1. In Section 3.2, we report results by marital status and education.

Sample. Our preferred sample consists of adults aged 21-59. An adult is said to be parent of a school-age child if one of their own children in the home is between the ages of 5 and 17. Households whose only children are under age five are excluded to isolate the impact of school-age children on labor supply. The age restriction on adults captures 98 percent of parents with school-age children. Thus, this restriction ensures that parents are compared to other adults of

[11] This element of equation (2) is shared by a simpler regression that maps hours worked to policy *within the sample of parents.* Since childless adults are excluded in the latter, identification rests entirely on within-area variation in policy. The key difference between these approaches is that equation (2) allows that changes in school policy may be endogenous to changes in the state of the local labor market.

similar age who are more likely to share the same baseline propensity to work. In Online Appendix C.2, we show that the response of parental labor supply implied by equation (1) is larger if the sample includes childless adults over age 59, consistent with results in Garcia and Cowan (2022). As we show, though, the parental labor supply response in this context reflects—and is exaggerated by—a common component in hours shared by *all* adults under age 59.

In addition, our full sample encompasses the broadest geographic coverage possible. We include all 478 local areas constructed from county, metro, and state identifiers in the CPS (see Section 1). Analogous results for the 280 counties disclosed in the CPS are reported in Online Appendix C.3. Estimates based on the latter, more restricted sample are somewhat smaller (and less precisely estimated) than those reported below.

Control variables. There are two distinct groups of regressors in $\mathbf{x}_{iat}$, each of which was advanced in Garcia and Cowan (2022). The first consists of demographic controls: age (and age squared); race; marital status; educational attainment; an indicator for rural, urban, or suburban location; the number of children (of all ages under 18); an indicator for the presence of under-five-year-old children; and indicators of Hispanic heritage, foreign birth, veteran status, and disability.[12]

The second group of regressors tracks the trajectory of the pandemic. These controls are the cumulative number of cases and deaths; the new monthly number of cases and deaths; and indicators for nonpharmaceutical interventions, such as Stay at Home orders. While we include this group for the sake of completeness, our estimates of ψ are essentially invariant to them. The reason is that these controls are common across adults with and without children and, as such, are differenced away in regression models of parents' relative hours worked (see equation (1)).

A third potential group of controls includes respondents' experience with an industry and occupation. A case could be made for these controls insofar as the composition of job types may be correlated with areas' pandemic policies. However, these controls are problematic because they are not reported in the CPS for most nonparticipants.[13] Thus, the absence of an industry (and occupation) affiliation means that the agent does not work. As a result, there is little variation left in labor supply to account for. Still, Online Appendix C.4 reviews results with these controls and shows that the impact of in-person shares is estimated to be even smaller than reported below.

3.1 Full sample

[12] The only controls here that are not present in Garcia and Cowan are the indicators for rural-urban-suburban status and for the presence of under-five-year-old children in the home.
[13] Industry and occupation are collected of nonparticipants in the Outgoing Rotation Groups (ORG) who report that they have worked in the past 12 months. Note that the ORGs *as a whole* make up only one quarter of the CPS sample.

We proceed to estimate the standard two-way fixed effects model in equation (1). Table 2 presents estimates for two outcomes: weekly hours worked and an indicator for employment. We also report results for two periods: the longer one spans all of 2020-21 but for the summer months, whereas the shorter period covers the 2020-21 school year (September 2020 – May 2021). Finally, for each period and each outcome, we report results separately for men and women.[14]

Consider first the results for the longer sample period that spans all of 2020-21. The main parameter of interest is ψ, which measures the response of parents' hours worked to school policy. Among women, a shift from fully virtual to fully in-person instruction implies an increase in hours worked of nearly 0.6 per week. It turns out that the response of total weekly hours is entirely accounted for by the extensive margin, as mothers' employment rate is estimated to increase by two percentage points.[15] Remarkably, the overall hours response among fathers is nearly identical, but reflects a shift along the intensive margin. As we shall see, the importance of the intensive margin to fathers is a robust result. Finally, it is noteworthy that β, which captures the market-wide hours response, is estimated to be positive and statistically significant, suggesting that school policy over this period may indeed stand in for broader shifts in the propensity to work.

Table 2: Estimates of Equation (1)

	Weekly hours		Employment	
	All 20-21	School 20-21	All 20-21	School 20-21
Coefficient		Women		
In-person share, β	1.192***	-0.881	0.019**	-0.013
	[0.338]	[0.595]	[0.008]	[0.014]
In-person × kids, ψ	0.582*	2.118***	0.020***	0.051***
	[0.304]	[0.590]	[0.007]	[0.014]
Number of obs.	447,899	228,550	447,899	228,550
Coefficient		Men		
In-person share, β	1.279***	-0.043	0.029***	0.016
	[0.382]	[0.649]	[0.008]	[0.014]
In-person × kids, ψ	0.567*	1.454***	-0.010	0.009
	[0.314]	[0.588]	[0.007]	[0.012]
Number of obs.	432,856	221,080	432,856	221,080

Note: Each column within each panel is a separate regression. In addition to the coefficients listed in the table, each regression includes the controls described in the main text (see "Control variables"). Standard errors are clustered at the geographic area level. "All 20-21" pools data for all of 2020 and 2021 but for the summer months (June, July, and August). "School 20-21" refers to the period September 2020 to May 2021. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

[14] Regressions in the main text are unweighted in view of arguments in Solon et al. (2015). For the sake of completeness, weighted results are included in Online Appendix C.6. On the whole, weighting makes little difference.
[15] Indeed, the higher employment rate implies a gain in total hours that slightly exceeds the estimated increase in weekly hours reported in the table, although the difference between the two is statistically insignificant.

Next, we turn to the 2020-21 school year. These results paint quite a different picture than the full 2020-21 sample. First, the overall hours response among parents is notably higher: a shift from fully virtual to fully in-person now implies an increase in mothers' relative labor input of just over two hours per week. Fathers' labor supply also appears to be more elastic, even if it is not quite as responsive as that of mothers. Once more, though, the margin of hours adjustment differs across men and women: the intensive margin dominates for fathers but is essentially unimportant for mothers. Meanwhile, the market-wide response to policy (β) is now insignificantly different from zero for both parents. We have confirmed that these differences across the two periods reflect the influence of the months that preceded the 2020-21 school year (e.g., January – May 2020) and *not* the months that followed (September – December 2021).

The parameter instability evident in Table 2 may reflect model mis-specification. One concern about equation (1) is that it omits controls for broader trends in parents' relative labor supply. For instance, if parents' jobs were generally less exposed to the initial turbulence of the pandemic, it would look as if their labor supply is somewhat insensitive to shifts in school policy that coincided with pandemic-related disruptions.[16] A corollary is that market-wide reactions to these disruptions may be large and will be reflected in a significant response to (correlated) changes in school policy. Notably, these concerns are likely less acute later in the pandemic. Thus, the absence of controls for such trends may lead to different estimates of ψ across different periods.

In view of this concern, we re-estimate the regression with additional controls for parent-specific trends in labor supply. Formally, this step merely requires the introduction of month-by-parental status fixed effects ($\theta_t \mathbb{k}_{it}$ in equation (2)). The first two columns of Table 3 report the results. To conserve space, we present only the response of hours here. In Online Appendix C.1, we confirm that the extensive margin continues to play an outsized role in women's labor supply response but matters little for men.

Under this specification, the adjustment of hours to in-person instruction is now remarkably stable across time. Among women, a shift from fully virtual to fully in-person instruction yields an increase in weekly hours of around 2.4—regardless of the sample period. The response among men is somewhat smaller—weekly hours increase by around 1.8—but again, is virtually unchanged across sample periods. Thus, as anticipated, the parameter instability in Table 2 reflected the failure to control for broader trends in parental labor supply. With the addition of these controls, the results for all periods are comparable to the results for the 2020-21 school year in Table 2.

[16] Lofton et al. (2021) document that, in the first few months of the pandemic, fathers experienced the least dramatic decline in employment and employed mothers experienced the smallest decline in weekly hours worked.

Table 3: Estimates of Equation (2)

	All 20-21	School 20-21	All 20-21	School 20-21
Coefficient		Women		
In-person share, β	0.553	-0.993*	1.405***	-0.118
	[0.394]	[0.601]	[0.406]	[0.670]
In-person × kids, ψ	2.360***	2.463***	-0.040	-0.101
	[0.634]	[0.633]	[0.672]	[1.129]
Number of obs.	447,899	228,550	447,899	228,550
Coefficient		Men		
In-person share, β	0.898**	-0.126	1.490***	0.964
	[0.400]	[0.654]	[0.427]	[0.757]
In-person × kids, ψ	1.888***	1.776***	-0.051	-1.737
	[0.645]	[0.628]	[0.706]	[1.194]
Number of obs.	432,856	221,080	432,856	221,080
Month × parent	Yes	Yes	Yes	Yes
Area × parent	No	No	Yes	Yes

Note: Each column within each panel is a separate regression. The dependent variable is the number of hours worked per week. See the notes of Table 2 for the other control variables included. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

Just as there may be parent-specific trends in hours worked, there may be parent-specific factors behind mean hours in a given area. These factors drive a wedge between the average hours of parents and childless adults within an area and may vary across areas. Such spatial differences pose a challenge to estimation if they reflect a "deep" feature of the local market but are nevertheless correlated with (average) 2020-21 in-person instruction rates. The reasons for any such correlation are perhaps not immediate, but it is easy all the same to add controls for spatial heterogeneity. As previewed in Section 2, we include area-by-parental status fixed effects ($\zeta_a \mathbb{k}_{it}$ in equation (2)), which control for mean differences across areas in relative parental labor supply.

The impact of these controls, shown in the final two columns of the table, is considerable: the response of parental labor supply to a change in the in-person share vanishes entirely. These results indicate that, once aggregate time trends are controlled for, the coefficient ψ is identified principally off cross-area comparisons of parents' relative hours worked. With additional controls for average regional differences in labor supply, the estimated effects of school policy disappear.

One could fairly question, though, if Table 3 "over-controls" for unobserved heterogeneity. With the addition of area-by-parental status terms, much of the important variation in school policy is now captured by other regressors. In other words, this perspective asserts that differences in average school policy across regions represent plausibly exogenous variation that may be, and indeed should be, used to recover the effect of in-person instruction.

This claim is at least partially testable. By this view, differences in average policies emerged for reasons unrelated to the "deep" attributes of a region that generally shape parental labor supply. In this case, there is no reason to expect average policies in the pandemic to predict *pre*-pandemic labor supply. This line of argument suggests a simple placebo test: are average in-person instruction shares in 2020-21 correlated with parents' relative hours worked before 2020?

In fact, this correlation between pandemic-era policy and pre-pandemic hours is plainly evident in the data. Figure 2 illustrates this point for a single (and the largest) demographic group, non-Hispanic White people. The x-axis shows the average in-person share in each of our local labor market areas over the pandemic period, 2020-21. The y-axis is based on pre-pandemic hours data from the CPS. Specifically, the y-axis shows the local-area average of parents' hours less the average of childless adults' hours over the five years prior to the pandemic, 2015-19. Remarkably, parents' relative hours worked in the pre-pandemic period happen to be higher in areas where schools chose more in-person instruction in 2020-21.

To pursue this point further, we apply equation (1) to test if pandemic-era policy predicts pre-pandemic hours in the full sample of adults (ages 21-59). Hours worked and all individual-level controls are taken from the CPS over 2015-19. The policy term, which was formerly measured by monthly data on 2020-21 in-person shares (p_{at}), is now the area-level mean of the latter and denoted by p_a. Note that since the policy term is fixed over time, the market-wide response to in-person shares, as measured by β, is not separately identified from the area-level fixed effects (χ_a). Therefore, we report only the response of parents' relative hours, e.g., the coefficient ψ on the term, $p_a \mathbb{k}_{it}$. A significant estimate of ψ indicates that average policies predict parents' relative labor input prior to the pandemic.[17]

The placebo test is implemented for men and women with two variants of mean in-person shares. We first calculate an area-level mean p_a based on all of 2020-21 (excluding summer months) and then consider an alternative using only the 2020-21 school year. The results are most striking for the former, especially for mothers. In areas that selected full-time in-person instruction, mothers' relative labor input *prior* to the pandemic was nearly 3.4 weekly hours greater than in areas with full-time virtual instruction. Among fathers, in-person instruction implied 1.4 more weekly hours of work, although the latter is only marginally significant (at the 10 percent level).

[17] We have confirmed that the results are virtually unaffected if we insert month-by-parental status effects as in equation (2). Of course, *area*-by-parental status effects are excluded so that we may measure the extent to which the latter variation would be captured by cross-area differences in average policy.

Figure 2: Pandemic School Formats and Pre-Pandemic Hours Worked

Note: This figure plots (on the y-axis) the difference in average pre-pandemic weekly hours between parents and childless adults against (on the x-axis) the average in-person share in the pandemic period. Each marker is a local labor market area, as described in the text (see Section 2). The left panel is based on hours data among non-Hispanic White women ages 21-59; the right panel refers to non-Hispanic White men (in the same age range). The pre-pandemic period spans 2015-19, whereas the pandemic period covers 2020-21. In each period, the summer months (June-August) are excluded. The line of best fit in the left panel (among women) has slope 6.404 (s.e. of 1.335), and the line of best fit in the right panel (among men) has slope 2.642 (s.e. of 1.097).

Notably, these figures are comparable to—or even exceed, in the case of mothers—estimates of hours responses in the pandemic period (see Table 3). When we compute mean in-person shares based on 2020-21 school year data, the estimate of ψ for mothers falls to two hours per week but remains strongly significant. The analogue for men declines to about one hour and is, again, marginally significant. These results are reported in Online Appendix B, where we confirm that the placebo test also fails with measures of the in-person share other than SafeGraph.

To understand the connection between pandemic-era school policy and pre-pandemic hours worked, it is helpful to first consider what, in general, may shape spatial dispersion in (pre-pandemic) parents' labor supply. Market work entails at least two costs that bear especially on parental labor supply and likely vary in the cross section. (Each of these factors is present in a model sketched later in Section 5.) The first is the cost of school-age childcare. The second is the utility cost of foregone time with children, which in turn reflects (i) the cost *per* unit time and (ii) the total time away from home. Although there is little reason to suspect (i) varies spatially, there is at least one factor behind (ii) that *does* do so, namely, the time it takes to commute to work.

We next show that commute times and childcare costs are correlated with (pandemic-era) school policy and (pre-pandemic) parental hours. Their connection to school policy runs, in part, through their association with local partisan affiliation. As we noted, in-person shares were highest where support for Donald Trump ran highest. At the same time, commutes are longer in populated metro areas where Trump's vote share was low. We also find that school-age childcare costs are higher in areas less supportive of President Trump. The latter result may partly reflect a difference in regulation: maximum child-to-staff ratios are lower in more politically liberal areas, which likely contributes to a higher cost of care. Online Appendix B reviews these results in detail.

In addition, the Appendix reports on the connection between these two outcomes—commute length and childcare costs—and parental labor supply. To be sure, these estimates are simple correlations. Still, a statistically significant correlation is instructive since it suggests that any *other* related outcome, such as the in-person share, is also likely to emerge as an apparent contributor to parents' labor supply. We find that longer commutes and higher childcare prices are indeed associated with lower *maternal* working time. Interestingly, though, they are essentially uncorrelated with *paternal* working time. In qualitative terms, these results echo more careful, causal analyses, which uncover a response among mothers but not fathers. Examples include Black et al. (2014) on commuting time and Mumford et al. (2020) on childcare prices.[18]

Taking stock of our findings, we conclude with the following observations. First, the failure of the placebo test raises alarms around specifications such as equation (1). We view the latter as likely to overstate the hours response because of its failure to address the endogeneity of school policy. Second, insofar as there is still a causal effect of school policy, it is also reasonable to worry that regressions with a full set of spatial controls, such as equation (2), are over-saturated and, as a result, will fail to detect the hours response. Third, the source of the endogeneity of in-person shares is not fully resolved. Among mothers, the close connection between hours worked, commute times, childcare prices, and in-person shares suggests that school policy is reflective of more general and fundamental forces at play in the local labor market. This narrative does not apply neatly to fathers, though.

For the remainder of this section, we present results based on both equations (1) and (2). For the reasons discussed, we see these results as likely to bound the true effect of school policy.

Sensitivity analysis, I: Policy measures. In addition to Parolin and Lee's (2021) analysis of SafeGraph data, there were other efforts to document the instruction modes at primary and secondary schools. We reexamine the impact of in-person instruction through the lens of two sources independently developed by the CSDH and Burbio.

[18] For a review of research on the connection between child care prices and maternal labor supply, see also Blau and Currie (2006). Paternal labor supply has received less attention in this context. As a complement to the empirical analysis in Mumford et al. (2020), Guner et al. (2020) study a calibrated lifecycle model that predicts household hours of work will be reallocated *away* from fathers *to* mothers when the price of child care falls.

The CSDH is based primarily on school-level reports of the predominant instruction mode. The reports were usually submitted monthly to state education agencies over the course of the 2020-21 school year.[19] In total, 35 states provided school-level data to CSDH. In another 11 states where school-level data was unavailable, agencies collected information at the school district-level. The 46 states for which CSDH provides data account for 2,800 of 3,100 U.S. counties and over 90 percent of U.S. student enrollment.

CSDH standardizes reports of instruction mode, categorizing them into one of three groups: in-person, virtual, or hybrid. We aggregate the school- and, where needed, district-level data to the county level. A score of one is assigned to a report of "in-person" instruction and a score of zero to "virtual" instruction. The on-site portion of "hybrid" instruction is never specified, though. For this reason, we draw from the 2021 National Assessment of Educational Progress (NAEP), which surveyed 5,000 schools as to the number of in-person days in their hybrid schedule. The in-person share implied by the NAEP represents the score (between 0 and 1) for hybrid instruction. A county's in-person share is the enrollment-weighted average of scores across the three formats.[20]

Although state agency data represent an official record of instruction format, they are not without noise. The categorical nature of the data necessarily involves a certain degree of judgment. For instance, a schedule with two days per week of on-site instruction only for grades K-2 may be understood as a "hybrid" format by one (primary) school but a "virtual" format by another.[21]

Therefore, we next turn to Burbio, whose estimates are developed from entirely different sources. Burbio's analysts follow district websites, local news reports, and social media to track the instruction format of a sample of school districts in their assigned area. As in CSDH, the format is categorized as in-person, virtual, or hybrid, and we again use NAEP data to quantify the in-person content of the hybrid schedule. Relative to CSDH, Burbio offers less geographic coverage: in total, Burbio follows districts in just under 600 U.S. counties.

Despite their differences, CSDH and Burbio share an important feature that distinguishes them from SafeGraph. Both CSDH and Burbio document in-person instruction offered by schools. In many of these districts, though, a virtual option was available to parents. Calarco et al. (2021) report that, in their survey of parents in late 2020, 75 percent of children had at least some access to in-person instruction, but less than 60 percent attended school on-site. The incomplete take-up of in-person instruction is captured by SafeGraph since it tracks foot traffic on school grounds.

---

[19] In 11 states, though, the reports were made to the U.S. Department of Agriculture as part of a program to reach students who were eligible for reduced-price meals but who did not attend school on-site.

[20] NAEP state-level tabulations are published for 37 states. We assign a state's in-person share to each county within the state. Estimates for the Census region are used if state data is unavailable. See Online Appendix C.5 for more.

[21] The hybrid format is a quantitatively important mode in the CSDH data: the hybrid share of instruction in each state is at least 20 percent and is as high as two thirds (North Carolina).

The presence of incomplete take-up implies that labor supply responses to CSDH and Burbio data will be diluted (relative to the effect of treatment on the treated). However, the role of take-up also poses a challenge because it is endogenous to labor supply, and the implied bias is clear: a parent who wants to work is more likely to enroll children in in-person instruction. For both reasons, estimates off SafeGraph data are likely to exceed those based on CSDH and Burbio data.

We proceed to estimate the hours worked response to CSDH and Burbio measures of in-person instruction.[22] We first re-estimate the standard two-way fixed effects model in equation (1) and then add further controls for unobserved heterogeneity. The sample period is the 2020-21 school year, as these are the months for which CSDH is available. We focus here on the main parameter of interest, namely, the parent's relative hours response, ψ. The estimates may be compared to SafeGraph-based results in Tables 2 and 3.

A few themes emerge from Table 4. Just as we saw above, labor supply responses are insignificant (and sometimes of the "wrong" sign) if all controls for unobserved heterogeneity are included. Across the other specifications reported in Table 4, the estimates are uniformly smaller than their counterparts based on SafeGraph data. Nevertheless, maternal labor supply responses based on CSDH and Burbio are often statistically significant (in these specifications), with a range centered around one hour per week. Paternal responses are now generally insignificant, though. Finally, the maternal-paternal differential is somewhat larger in CSDH vis à vis Burbio data.

Table 4: Estimates with Alternative Measures of School Formats

	CSDH			Burbio		
Coefficient			Women			
In-person × kids, ψ	1.013**	1.423***	-1.158	0.608	1.011**	0.373
	[0.484]	[0.548]	[0.942]	[0.418]	[0.490]	[0.704]
Number of obs.	211,156	211,156	211,156	211,777	211,777	211,777
Coefficient			Men			
In-person × kids, ψ	0.211	0.448	-1.299	0.500	0.774	0.708
	[0.484]	[0.542]	[0.822]	[0.438]	[0.493]	[0.732]
Number of obs.	204,090	204,090	204,090	205,039	205,039	205,039
Month × parent	No	Yes	Yes	No	Yes	Yes
Area × parent	No	No	Yes	No	No	Yes

Note: The controls in each column are identical to those used in Tables 2 and 3 but for the measurement of the in-person share, which is now drawn from CSDH or Burbio. For each of the latter, one column reports estimates of equation (1), and the remainder of the columns include some combination of parental status controls in equation (2). *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

[22] Again, results characterizing the extensive margin are provided in Online Appendix C.1.

Specifically, the maternal labor supply response exceeds its paternal counterpart by as much as one hour per week given CSDH data; the gap is one quarter of an hour in Burbio data. The differential lies in between the results in SafeGraph data.[23]

Sensitivity analysis, II: Sample of parents. Thus far, we have defined school-age children as kids between ages 5 and 17. However, children whose age is near the top of this range may not require much parental supervision (see Blau and Currie, 2006). This observation leads us to examine the extent to which parents' labor supply responses vary based on their children's ages.

We present results for three age ranges. To start, we include parents in the sample only if their eldest child is between ages 5 and 13, excluding children of high school age. Next, we narrow the age range to include only parents whose eldest child is between ages 5 and 9. Finally, to put these estimates in context, we consider a sample that includes parents only if their eldest child is *older* than 13 (but less than 18). For the sake of brevity, we again focus on the main coefficient of interest, ψ, and report results only for the full sample period, 2020-21 (but with the summer months omitted). The in-person share is from SafeGraph.[24]

Table 5: Estimates with Alternative Definitions of School-Age Children

	Age range of school-age children					
	14−17	5−13	5−9	14−17	5−13	5−9
Coefficient	Women					
In-person × kids, ψ	1.453	2.439**	2.667**	0.041	0.114	-0.504
	[0.960]	[0.711]	[0.863]	[1.066]	[0.757]	[0.876]
Number of obs.	326,585	406,570	359,375	326,585	406,570	359,375
Coefficient	Men					
In-person × kids, ψ	1.001	2.057**	2.944**	-0.696	0.017	0.788
	[0.985]	[0.686]	[0.823]	[1.232]	[0.806]	[1.011]
Number of obs.	334,121	400,858	363,400	334,121	400,858	363,400
Month × parent	Yes	Yes	Yes	Yes	Yes	Yes
Area × parent	No	No	No	Yes	Yes	Yes

Note: This table presents estimates based on equation (2) under different definitions of "school age". Each column's sample includes childless adults and parents whose eldest child's age lies within the range in the column header. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

[23] The larger responses in SafeGraph data were arguably anticipated by our discussion of take-up. Alternatively, the coarseness of the classifications in the CSDH and Burbio data may lead to a form of attenuation bias. We explore this point further in Online Appendix C.5.

[24] While the use of this measure ensures a certain consistency with earlier specifications, it reflects on-site activity across all grades. Online Appendix C.5 confirms that regressions with in-person shares (again, from SafeGraph) tailored to the age ranges of the children in each sample yield similar conclusions.

Results are presented in Table 5. Consider first the estimates in the three far-left columns. Among parents of children ages 5-9, a shift from fully virtual to fully in-person implies a gain in weekly hours of between 2.5 and 3 for mothers and fathers. When we extend the upper limit of the age range to 13, the response of fathers' hours falls notably, but the decline among mothers is more muted. Finally, in households with older children (ages 14-17), the hours response of fathers is halved further and is insignificant for both parents. Clearly, the significant response among all parents of children ages 5-17 in Table 3 largely reflects the behavior of parents of younger children. The results on the right panel show, though, that the inclusion of area-by-parental status effects again eliminates the significance of the estimates.

3.2 Education and marital status

In line with related research, we next ask if parental labor supply responses to virtual instruction (that is, ψ) differed by marital status and/or educational attainment. The analysis will focus on the response of total weekly hours. Online Appendix C.1 reviews results for employment. In addition, we retain throughout controls for parent-specific trends (time × parental status effects) but exclude parent-specific spatial controls (area × parental status effects). We have confirmed that the inclusion of the spatial controls eliminates the statistical significance of the estimates, just as they do in Section 3. One might then view the results below as the strongest case that one could present for a role of school policy in parental labor supply.

Education. We first consider the role of education in the labor supply response to school policy. We divide our sample into a noncollege group—workers with less than a four-year degree—and workers who completed college. We further split each of these two groups by gender. Results are reported in Table 6.

Consider first the estimates for women in the top panel of the table. Among the noncollege educated, a shift from fully virtual to fully in-person implies an increase in weekly hours of just over two. The response among college graduates is only slightly smaller; the two responses are not statistically distinguishable from one another. Thus, among women, college experience is not a strong predictor of the labor supply response to school policy.

Table 6: Estimates by Educational Background

	Noncollege	College	Noncollege	College
Coefficient	Women		Men	
In-person × kids, ψ	2.074***	1.818*	1.999***	1.078
	[0.771]	[1.001]	[0.751]	[0.863]
Number of obs.	266,258	181,641	284,723	148,133

Note: Each column header reports the sample used in the regression. A college (noncollege) graduate is one who did (not) complete a four-year degree. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

The education gradient among men is more evident. College-educated men do not significantly adjust their hours worked in response to variation in school policy. By contrast, among noncollege men, a shift from fully virtual to in-person implies an increase in weekly hours that is comparable to the response observed among (noncollege and college-educated) women.

A corollary of these results is that male and female labor supply *within* the college group diverged. This point is sharpened if we consider households with two college-educated spouses. The hours response among mothers in these households is somewhat larger than the response among college graduate mothers as a whole. A shift from fully virtual to in-person now implies an increase of 2.8 hours per week (see Online Appendix C.7). By contrast, fathers' labor supply in such households is slightly less responsive than shown in Table 6 and statistically indistinguishable from zero. This imbalance between spouses is, again, evident only among the college educated. In households where both spouses are noncollege educated, hours responses are almost identical.[25]

Marital status. Marital status is often singled out as a critical factor in pandemic-era labor supply decisions, but the direction of its effect is not obvious *a priori*. One view suggests that married partners can share childcare and work responsibilities, reducing the need for either parent to exit employment to supervise children. By contrast, a single parent may be the only adult who could supply supervision. Thus, even if one worker in a married household reduces labor supply, we may expect labor input among married workers on average to decline by less in the event of virtual instruction (Garcia and Cowan, 2022). Another view highlights that a single parent who is the lone breadwinner in the household may have a strong incentive to arrange some other form of childcare to carry on in the workforce. Under this view, the propensity of a single parent to work could exceed that of any married individual (Heggeness and Suri, 2021).

This debate over the role of marital status hints at another, related consideration: the broader composition of the household. The reference to a single parent, as above, often assumes the parent is the lone adult in the household. However, a single parent can receive childcare support within a co-residential arrangement with other adults. In this situation, the labor supply of single parents may still be relatively insulated from variation in school policy.

This consideration is likely to be empirically relevant. In the CPS, almost 55 percent of unmarried mothers live with at least one other adult. Among the latter, the most common arrangement is cohabitation: over 40 percent of these unmarried mothers live with an unmarried partner. In addition, just over one quarter of them live with an older parent; just under one quarter live with an adult child (aged 18 or older); and over 5 percent live with a brother or sister.

[25] Online Appendix C.7 asks if these patterns reflect intra-household differences in earnings opportunities. Intuitively, if a college-educated father is more likely (than a noncollege graduate) to have earnings exceeding those of his spouse, he may be more likely to continue working during a school closure. In fact, this premise finds little support in the data.

Our estimates in Table 7 confirm that household composition mediates the role of marital status. For women, the response of hours worked among unmarried mothers as a whole is similar to that of the married sample: a shift from (fully) virtual to in-person implies a gain between 2.2 to 2.6 weekly hours. However, the estimate for the unmarried masks a significant difference between mothers with and without other adults in the household. Among lone-adult mothers, hours worked are notably more responsive: a shift from (fully) virtual to in-person implies an increase of almost 4.3 weekly hours. By contrast, the response of unmarried women in co-residential arrangements (not shown) is 1.4 hours per week and statistically insignificant.

The narrative for men is rather similar. For instance, as we saw for women, the results differ little across marital status. The estimate for married fathers does come in below its counterpart for mothers, but one could not reject equivalence. The estimate for the unmarried is relatively imprecise, in no small part because few unmarried men live with their children.[26] In light of the latter result, the estimate of ψ for lone-adult men is significant and almost identical to its counterpart for lone-adult women.

Table 7: Estimates by Marital Status

	Married	Unmarried	Lone adults
Coefficient		Women	
In-person × kids, ψ	2.256***	2.591**	4.275***
	[0.788]	[1.032]	[1.178]
Number of obs.	242,743	205,156	67,592
Coefficient		Men	
In-person × kids, ψ	1.824***	1.657	4.372*
	[0.661]	[1.455]	[2.248]
Number of obs.	223,471	209,385	61,954

Note: The column header reports the composition of the sample used in the regression. A "lone adult" is a respondent who does not live with any other individual age 18 or over. The sample is 2020-21 with summer months excluded. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

It is instructive to compare these results with other research in this area. Our estimates for men are similar to those in Garcia and Cowan (2022), although the latter's estimate for the unmarried is statistically significant. The estimates for women are higher, especially for married mothers. Online Appendix C.6 traces these discrepancies to differences in sample and controls. Our estimates also tend to exceed those reported by Hansen et al. (2022), who find no labor supply

[26] Whereas 22.5 percent of unmarried mothers live with their kids, just 9 percent of unmarried fathers do so.

response by unmarried mothers or fathers.[27] In short, the estimates are at, or near, the top end of the range in this literature. In Section 5, though, we argue that even these results are, in a sense made precise, unexpectedly small.

4. Estimates from Time Use Data

Our analysis of CPS data suggests that a shift from a virtual to an in-person format was associated with an increase of no more than two to four weekly hours of work. The suspension of on-site instruction, however, removed over 30 hours of school-provided supervision. Thus, the labor supply response suggests that parents must have adjusted to school closures on other margins.

To examine time use adjustments more broadly, we turn to the American Time Use Survey (ATUS) (Hofferth et al., 2020). Our ATUS sample is selected to conform to the extent possible with our treatment of the CPS. We again restrict attention to individuals ages 21-59 who are childless adults or parents of school-age children. (As above, parents are excluded if all their children are under school age.) The ATUS sample excludes data from mid-March to mid-May 2020, as the pandemic forced the suspension of field work from mid-March to mid-May 2020.

For each respondent, we observe a minute-by-minute diary of a single day that describes how, where, and with whom they spent their day. The days of the week are not uniformly represented, though. Rather, half of respondents report activities for Saturday or Sunday. Given our interest in how time use responds to the status of school instruction, our baseline sample restricts attention to school days, Monday to Friday. This restriction leaves us with 3,278 individual observations. As an alternative, we use the full sample of days but apply ATUS weights, which adjust for the oversample of weekend days. We focus on our baseline in the main text and report weighted OLS results in Online Appendix D.

The ATUS diary data enable us to examine if the return of in-person instruction affected the distribution of time use across a range of activities. The respondent's diary entries are assigned one of many detailed activity codes. We group these activities into a few broad categories: work, leisure, home production, childcare, commute time, and sleep. We then estimate how hours spent in each category responds to variation in instruction format. As in Section 3.2, the specification follows equation (1) but with parental-status-by-time effects. In addition, we include a fixed effect for each day of the week. Since the data are daily, the point estimates are scaled to express them on a weekly basis and, therefore, comparable to estimates from the CPS.

[27] Hansen et al. results for the unmarried also diverge from those in Garcia and Cowan. The reasons for these differences have been difficult to identify. One possibility is that the in-person shares are different; Hansen et al. produce their own with SafeGraph data whereas we and Garcia and Cowan use Parolin and Lee (2021).

Remarkably, we find little evidence of adjustments in parents' time use. Specifically, the reinstatement of in-person instruction has, on the whole, no significant impact on any major category of time use, from work to leisure, childcare, and home production. These results obtain when we split the sample by college attainment. See Online Appendix D for details. What we take from this exercise is that, whatever are the "true" effects of school policy, they are not large enough to detect in the modestly sized ATUS sample. Perhaps this result should not be a surprise: the much larger sample in the CPS yields a mean response of market time of just two weekly hours. [28]

However, there is a sense in which these regressions do not leverage the richness of the ATUS. In addition to the activities undertaken, the ATUS records *where* each activity takes place. The ATUS also asks if there was a child in the respondent's care. Thus, we can observe if a parent is supervising a school-age child while working at home. This level of detail enables us to move beyond standard labor supply analyses by recognizing that market work was not necessarily performed in isolation from other activities but may be "bundled" with them.

Table 8 reports on the role of working from home as a means of both supplying childcare and market time. To start, the far-left column reiterates that total working hours in the ATUS are insensitive to instruction format. The next column reports results for total hours working at home. Interestingly, this, too, does not respond significantly. However, the final column shows that time spent working at home while caring for children is responsive to instruction format, but only among college graduates. After a shift from fully virtual to in-person instruction, college-educated parents reduced time in this activity by 6 hours per week. Thus, college graduates continued teleworking after in-person instruction resumed but no longer supervised children while doing so. Online Appendix D shows that this result stems to a large degree specifically from college educated mothers, but standard errors in these subsamples are rather large (which is why we pool men and women in Table 8). The response of the noncollege group is smaller and statistically insignificant, consistent with evidence that this group had fewer telework opportunities (Mongey et al., 2020).

These results strongly suggest that college educated parents relied on telework to sustain their hours worked when instruction was virtual. This inference should be qualified in two respects, though. First, among college graduates, the increase in telework while engaged in childcare (6 weekly hours) still represents a portion of the time that had to be filled when in-person instruction was suspended. Second, the non-college educated in the ATUS also smoothed hours worked but

---

[28] The discrepancy between CPS and ATUS-based results is unlikely to reflect systematically different measurements of hours worked. Research has found substantial agreement between the two sources (Frazis and Stewart, 2004, 2014).

did not rely on telework. Thus, if telework had been infeasible, it is unlikely that the labor supply of the college educated would have fallen 6 weekly hours. Rather, the college group would have

Table 8: Work at Home, Childcare, and Instruction Format

	Work	Work at home	Work at home + childcare
Coefficient		All	
In-person share, β	0.565	-3.300	1.553
	[3.973]	[2.960]	[0.935]
In-person × kids, ψ	1.839	-2.685	-5.897***
	[4.256]	[3.952]	[1.535]
Number of obs.	3,278	3,278	3,278
		Noncollege	
In-person share, β	-1.657	-1.825	-0.455
	[6.338]	[3.170]	[0.703]
In-person × kids, ψ	1.125	1.848	-1.903
	[7.168]	[4.547]	[2.290]
Number of obs.	1,623	1,623	1,623
		College	
In-person share, β	6.540	-2.474	2.800
	[5.37]	[4.825]	[1.566]
In-person × kids, ψ	0.315	0.990	-6.652**
	[6.146]	[7.168]	[2.704]
Number of obs.	1,561	1,561	1,561

Note: Each column by panel is a separate regression, with the implied number of hours per week spent in each activity as the dependent variable. (Daily hours are multiplied by five.) The sample consists of observations between Monday and Friday. Relative to equation (1), we also include fixed effects for days of the week and parent status×month. Standard errors are clustered at the geographic area level. "Work at home" is the number of work hours carried out in one's own home or another home. "Work at home + childcare" measures the number of hours where "work at home" is the primary activity and "childcare" is the secondary activity. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

taken up (to some degree) measures adopted by the noncollege group to cope with shifts in instruction format.[29]

Insofar as telework may not fully fill the gap (in time) left by the suspension of in-person instruction, we turn our attention to another margin of adjustment: the utilization of nonparental childcare. A survey fielded in late 2020 by Calarco et al. (2021), and analyzed further in Yang et

[29] To more credibly recover causal effects, one could try to exploit plausibly exogenous variation in workers' access, or "exposure", to telework. However, measures of access are based on occupations (Dingel and Neiman, 2020) and are not easily mapped to nonemployed survey respondents.

al. (2022), reports specifically on the use of non-center-based, or *informal,* care, which includes unpaid care by family and friends as well as in-home paid care (e.g., nannies). Informal care was reported by 60 percent of surveyed families, three quarters of whom had at least one school-age child. Parents with school-age children were said to utilize such care if there were providers who "helped with learning at home" in Fall 2020. Thus, we read these data to suggest a promising role for nonparental supervision in periods of virtual instruction.[30]

The ATUS also allows us to examine a role for nonparental care, albeit in a more limited form. For each adult aged 60 years or older, we calculate the number of hours per week spent with children under age 18 that are *not* the respondent's son or daughter. This sum excludes time spent at work to try to identify unpaid, informal care of the sort that a grandparent or other older relative might provide. Table 9 reports how these hours of care vary with the in-person share of instruction. Note that the specification in this context *excludes* the parental status indicator; the covariate of interest is simply the in-person share. Hence, the identification assumption is that school policy did not systematically vary with older respondents' preferences for caregiving.

The estimates from ATUS indicate that older respondent's caregiving lessened when in-person instruction resumed. In the full sample, the resumption of fully in-person instruction is estimated to reduce older respondents' time with children by nearly 3.4 hours per week. The latter result in turn stems in large part from the behavior of older women, whose weekly hours of caregiving decline by nearly 3.9 when in-person instruction returns. We treat these estimates as further suggestive evidence that non-parent adults assumed more childcare responsibilities when in-person instruction was suspended. Online Appendix D shows that these results do not vary much by the education of the older respondent. However, in the weighted regressions (where ATUS sample weights are used), the response is more concentrated among noncollege educated women.

Table 9: Time with Others' Children and Local School Formats

Coefficient	All	Men	Women
In-person share, β	-3.385**	-2.547	-3.878*
	[1.480]	[1.734]	[2.165]
Number of obs.	2,425	976	1,354

Note: Each column is a separate regression, with the implied number of hours per week spent with other's children as the dependent variable. (Daily hours are multiplied by five.) Time spent with other's children includes all time spent with persons under 18 years old outside of market work. The sample includes individuals who are 60 years or older. Standard errors are clustered at the geographic area level. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

[30] Yang et al. (2022) finds that more highly educated parents used informal care more intensively, although the differences are statistically different only for those with advanced degrees. This result underscores our view that, even with the advantage afforded by telework, many of the college educated also turned to nonparental care.

5. Discussion

We now use a series of time allocation models to guide a discussion of our regression results. Our first aim is to formalize why estimates of the labor supply responses to school closures are unexpectedly small. We demonstrate this point within a simple environment where a single parent faces a one-for-one tradeoff (in time) between labor supply and childcare. Under any reasonable parameterization, the model implies labor supply responses that far exceed any reported estimate. We illustrate formally how telework can relax the work-childcare tradeoff and, therefore, mute the response of hours worked. At the same time, hours worked responses were modest even for populations, such as the noncollege educated, who were less likely to access telework. As noted above, this observation leads us to also consider a role for nonparental care, which enables parents to smooth their labor supply and ensure the provision of childcare. To illustrate this point, and to conclude the section, we show how to parameterize a model with nonparental care to yield parental labor supply behavior in line with our regression results.

A simple baseline. A single parent maximizes utility over consumption, leisure, and child development subject to two constraints on his or her time. The first constraint is that the allocation of her time across child supervision, leisure activity, and market work must add up to the total time endowment (normalized to 1). The second constraint is that the child must be supervised at all times.

To start, we assume there are only two forms of child supervision. There is a publicly provided form of supervision, which the parent takes as given. The notion of publicly provided supervision is a crude description of in-class instruction time, but it arguably the captures the dimension of in-person instruction that is most relevant to a study of parental labor supply. If the child is not in school, we assume he/she must be under the parent's full-time supervision. We introduce private nonparental care below.

To proceed, consider the parent's time constraints. Leisure is denoted by l; time allocated to child supervision by m; and market hours of work by n. Finally, we let g be time spent under publicly provided supervision. The time constraints specify that a parent's allocations add up to 1,

$$l + m + n = 1, \tag{4}$$

and that the child must be under school or parental supervision,

$$g + m = 1. \tag{5}$$

Together, equations (4) and (5) imply

$$l = g - n. \tag{6}$$

Thus, a decrease in on-site instruction time, g, lowers leisure one for one unless market hours fall.

Importantly, we assume g is taken as given by the parent. This approach rules out substitution from a school closed to in-person instruction to one that is open to it. Where this did occur, it often took the form of a movement from public to private school. However, the rates of reallocation to private institutions would seem to be far too small to make a material difference to our analysis (see Dee et al., 2022; Musaddiq et al., 2022).

As in Berlinski et al. (2023), we assume that period utility is given by

$$u(c, l, q) = \alpha \ln c + \beta \ln l + (1 - \alpha - \beta) \ln q, \tag{7}$$

where $\alpha, \beta \in (0,1)$. Period utility depends on market consumption, leisure, and a term, q, that indexes child development and is "produced" with both forms of supervision, g and m. From equation (5), though, we know that g implies m, which means q is pinned down by g: $q(g, m) = q(g, 1 - g)$. Since g is taken as given by the parent, q acts merely as an exogenous intercept in the utility function. A more substantive choice problem for q emerges if there is another source of supervision, e.g., a form of private care outside the home. We return to this matter later.

In what follows, we consider the model's predictions for labor supply. Notably, the model only yields interior solutions for n, although the observed hours responses among parents reflects movements on the extensive margin. In our view, what we sacrifice in realism in this regard is worth the tractability that it purchases. Moreover, our calibration will take care to account for evidence on the Frisch elasticity of total hours, which reflects both intensive- and extensive-margin adjustments.

Initial comparative statics. The first order condition for leisure implies

$$g - n = \frac{\beta}{\lambda w}, \tag{8}$$

where λ is the marginal value of wealth and w is the wage rate. Now suppose for the moment that the parent can borrow and lend sufficiently to smooth the marginal utility of consumption. Under this assumption, the right side of equation (8) is (practically) invariant to g. It follows that n moves one for one with a change in g. Intuitively, the demand for leisure does not change since its opportunity cost is fixed given λ and w. Therefore, n must adjust to offset a shift in g, which leaves l unchanged.

One might object to our interpretation of equation (8) on the grounds that such perfect (self)-insurance is unrealistic. While the transfers made available under the CARES Act (and, later, the American Rescue Plan) likely enabled households to smooth their marginal utility to a considerable extent (Wu et al., 2022), we nevertheless, consider an alternative to the case of perfect insurance.

In this alternative scenario, parents live "hand to mouth." Therefore, consumption must satisfy $c = wn$. It follows that $\lambda = \alpha/c = \alpha/wn$, and equation (8) becomes

$$l = g - n = \frac{\beta}{\alpha} n. \tag{9}$$

A perturbation to g yields a change in hours work equal to

$$\mathrm{d}n = \frac{1}{1 + \beta/\alpha} \mathrm{d}g. \tag{10}$$

Equation (9) says that β/α is identified by the ratio of leisure to market work time. Based on data from the American Time Use Survey, we then find that $\beta/\alpha = 1.1$. Therefore, equation (10) implies that an hour more of in-person instruction yields an increase in market work of approximately 0.5 hours.[31]

This prediction would seem to far exceed any of our empirical estimates (or of estimates elsewhere in the literature). A shift from fully virtual to fully in-person will reinstate, on average, 32.5 hours per week of on-site instruction (Planty et al., 2008). The result in equation (10) predicts that such a shift should lift labor supply by 16 hours per week. Conversely, the closure of schools to on-site instruction is predicted to dramatically reduce parental hours worked, which is arguably consistent with some predictions at the start of the pandemic (see discussion in Goldin, 2022). By contrast, our OLS estimates suggest a labor supply response no higher than 2-4 hours per week.

In fact, reasonable alternative calibrations of equation (10) yield even larger predicted changes in market work. To see why, it is straightforward to show that β/α is in fact the Frisch elasticity in this model. If we were to calibrate the latter based on standard estimates of the Frisch elasticity (rather than based on time-use data), the pass through from g to n would be larger. For example, consider a Frisch elasticity of 0.8 based on the review of evidence on intensive- and extensive-margin adjustments in Chetty et al. (2011). It follows that a shift from fully virtual to fully in-person instruction now yields an increase in market work of 18 hours.

Telework. A key assumption embedded in equation (7) is that parents cannot simultaneously perform market work while they supervise children. However, there is evidence from the ATUS

[31] This calculation excludes hours of sleep from leisure time. Alternatively, if we treat hours of sleep in excess of some minimum—say, 6 per night—as leisure, then $\mathrm{d}n/\mathrm{d}g = 1/3$, and the increase in market work implied by a shift to in-person instruction falls to from 16 hours (see below) to just under 11.

suggesting that telework opportunities, at least for college graduates, enabled parents to provide some childcare even as they continued to supply labor. We illustrate a tractable way to capture this notion of telework in the model.

The new ingredient is a time aggregator function. The idea behind this function is that a parent may supply 8 hours of market work and 2 hours of childcare in under 10 hours. That is, the two activities may, to some degree, be done concurrently. Formally, the time aggregator function maps time engaged in market work, n, and time engaged in childcare, m, into the *total* time that has passed while engaged in one or both activities. The function has the form,

$$t(m,n) = (m^{\rho} + n^{\rho})^{1/\rho}, \tag{11}$$

where $\rho \geq 1$. The time constraint (7) then generalizes to $l + t(m,n) = 1$. Leisure, l, is defined as the absence of any other activity and, therefore, enters the time constraint separably (outside of t).

Equation (11) encompasses two polar cases. The first is $\rho = 1$, which recovers the original time constraint (7), $l + m + n = 1$. This case corresponds to the standard assumption that two activities are perfectly rivalrous—an hour of market work is done to the exclusion of an hour of childcare. The second is the limit where $\rho \to +\infty$, which implies that $t(m,n) \to \max\{m,n\}$. In this case, the two activities are perfectly *non*rivalrous. To illustrate, if $m > n$, an increase in market work can be completed within the time already allocated to childcare. More generally, the activities can be performed concurrently up to (of course) the minimum of the two.

These two polar cases are bridged by a continuum of finite $\rho > 1$. In this interior region, a few properties of equation (11) will be important. First, equation (11) implies $t_m \equiv \partial t/\partial m \in (0,1)$ and, similarly for market work, $t_n \equiv \partial t/\partial n \in (0,1)$. In words, another hour of any activity absorbs less than an hour of new time, because some portion of it is done concurrently with the other activity.[32] Therefore, we say the *time price* of an activity is less than one. Second, the time price of an activity increases in the time allocated to it (e.g., t is convex) and decreases in the time allocated to the other activity (e.g., $\partial^2 t/\partial n \partial m = \partial^2 t/\partial m \partial n < 0$). The intuition here is that, if m is large relative to n, a parent can identify many (new) childcare tasks that can be done concurrently with market work but few (new) work tasks that can be done jointly with childcare. Therefore, the time price of another hour of work is relatively small, but the price of another hour of care is relatively high.

[32] In the limit $\rho \to +\infty$, these derivatives are zero or one. Intuitively, if $m > n$, any market work can be done with current childcare, which implies $t_n = 0$. Conversely, if m rises, there is no scope to multi-task further, e.g., to complete a new childcare task jointly with current market work. Therefore, $t_m = 1$.

These properties formalize the sense in which equation (11) yields a motive to "multi-task." Since the time price of market work falls as childcare time rises, the parent has a strong incentive to elevate hours worked, too. This motive to multi-task is strengthened at higher values of ρ. To see this point, note that the time price of another hour of market work relative to childcare is given by $t_n/t_m = (m/n)^{1-\rho}$. At higher values of ρ, the relative price of market work falls more rapidly as m grows relative to n, which means that parents will want n to partially "catch up" to m. In the limit where $\rho \to +\infty$, the catch-up is complete, e.g., a (momentary) increase in m above n leads to an equal increase in n.

Consider now the choice of n. The first-order condition is

$$l = 1 - t(m,n) = \frac{\beta}{\lambda w} \cdot \frac{\partial t}{\partial n}. \tag{12}$$

A decline in publicly provided supervision, g, now has two effects. The first is familiar and relates to the left side of equation (12): $m = 1 - g$ must rise, which would diminish leisure all else equal. To stem the fall in leisure, labor supply is reduced.[33] The second effect is novel and operates via the time price of market work, $\partial t/\partial n$, on the right of equation (12). An increase in m reduces this price since $\partial^2 t/\partial n \partial m < 0$. As a result, there is now a motive to supply *more* labor when g falls, mitigating the fall in market time due to the first effect.

To investigate these points more fully, we inspect the comparative static of n with respect to g. Under perfect insurance ($\mathrm{d}\lambda = 0$), we have

$$\frac{\mathrm{d}n}{\mathrm{d}g} = \frac{1 - (\rho - 1)/\phi(l)}{(n/m)^\rho + (\rho - 1)/\phi(l)} \cdot \frac{n}{m}, \tag{13}$$

where $\phi(l) \equiv (1 - l)/l$ and $m = 1 - g$. When $\rho = 1$, equation (13) collapses to $\mathrm{d}n/\mathrm{d}g = 1$: market work is reduced one for one with a fall in g. Values of $\rho > 1$ can attenuate the decline in labor supply, e.g., $\mathrm{d}n/\mathrm{d}g \in [0,1)$. In fact, equation (13) implies that the decline is eliminated entirely if $\rho = 1 + \phi(l)$. The term $\phi(l)$ signals the degree of curvature over l in the utility function: if $\phi(l)$is large, the marginal utility of leisure rises rapidly with any reduction in l, which prompts the parent to reduce market hours instead. Thus, to the induce $\mathrm{d}n = 0$, the motive to multi-task must be strong enough to overwhelm the force of this curvature. Although arguably knife-edge, this case is instructive because the regression results (Section 4 in particular) suggest that

[33] However, the quantitative impact of even this "conventional" channel is different under equation (11). To illustrate, suppose m rises by one, but we wish to prevent any fall in leisure. When $\rho = 1$, n also declines by one. However, for $\rho > 1$, the required decline is shaped by the form of t and equal to $\mathrm{d}n = t_m/t_n$.

telework may have insulated the labor supply of some parents when in-person instruction was suspended. The technology in equation (11) is flexible enough to reproduce this outcome.

Nonparental care. Thus far, we have assumed that a child must be supervised by her school or parent. However, two data points strongly suggest that time in private nonparental care was an important margin of adjustment to school closures. First, labor supply responses are relatively modest even among workers with little access to telework (i.e., the noncollege educated). Second, the stasis in parents' time use more generally—the absence of a clear response of any other activity—is hard to rationalize unless the loss of in-person instruction time was filled by nonparental care. Therefore, we incorporate into the model the choice of time in private nonparental care, denoted here by x. The analogue to equation (5) is then

$$g + m + x = 1, \tag{14}$$

which says that a child is supervised by a school, parent, or private third party.

The revised time constraint implies a simple, but potentially substantive, change in labor supply. The first order condition for hours worked is now

$$n = g + x - \frac{\beta}{\lambda w}. \tag{15}$$

Market work, n, no longer necessarily moves one for one with in-person instruction time, g. Rather, market work moves one for one with the sum of time outside of parental care, $g + x$. Therefore, if private nonparental care (x) rises to offset a decline in publicly provided supervision (g), the labor supply response will be muted.

Each form of supervision is an input into the child's development. A particularly tractable specification for the development "production" function is given by

$$q = g^{\gamma} q_{\!,}(m, x)^{1-\gamma}\,, \ \text{with} \tag{16}$$
$$q_{\!,}(m, x) = \left(\mu^{1-\psi} m^{\psi} + (1-\mu)^{1-\psi} x^{\psi}\right)^{1/\psi}$$

and where $\gamma \in (0,1)$ and $\psi \leq 1$. Equation (16) uses a Cobb-Douglas outer nest to aggregate on-site instruction time (g) and a "bundle" of private care ($q_{\!,}$). The latter inner bundle is a CES aggregate with elasticity of substitution between parental (m) and private nonparental care (x) given by $(1-\psi)^{-1}$. The CES form is arguably the most popular specification since at least Cunha et al. (2010), though Del Boca et al. (2014) consider a Cobb-Douglas form. The Cobb-Douglas outer nest renders g and $q_{\!,}$ logarithmically separable, which simplifies the algebra to follow.

The optimal choice of any form of care trades off the value of another hour of time to the child with the price of that care. The price of parental care is the foregone market wage, w, whereas nonparental care has price per unit time denoted by p. It follows that, at the optimum, the marginal product of parental care net of nonparental care is equated to its relative price, $w - p$. Once the time constraint $m = 1 - g - x$ is then imposed, we recover the optimal choice of x, and therefore n. Note that p is "small" here in the sense that $w > p$ to account for the prevalence of informal care, e.g., supervision by friends, neighbors, grandparents, or older children (Yang et al., 2022).

We may now consider how parental labor supply responds to a shift in publicly provided supervision, g. As shown in Online Appendix E, this comparative static may be written as

$$\frac{\mathrm{d}n}{\mathrm{d}g} = \frac{z(\xi;\psi)}{1 + z(\xi;\psi)}, \tag{17}$$

where $\xi \equiv x/m = x/(1 - g - x)$ is nonparental time per hour of parental care and

$$z(\xi;\psi) \equiv \frac{\left(\frac{\mu\xi}{1-\mu}\right)^{\psi-1} + (1-\psi)\xi^{-1} - \psi}{\left(\frac{\mu\xi}{1-\mu}\right)^{1-\psi} + (1-\psi)\xi - \psi}.$$

The Appendix also establishes a few properties of z. First, $1 + z > 0$ given any initial optimum ξ consistent with $w > p$. Second, z declines in ψ and crosses zero at a $\hat{\psi} \in (0,1)$.[34] The comparative static displays the same monotonicity such that $0 \leq \mathrm{d}n/\mathrm{d}g$ for all $\psi \leq \hat{\psi}$ and $\mathrm{d}n/\mathrm{d}g < 0$ otherwise.

Thus, the choice of ψ enables the model to reproduce a range of market time responses. Where parental and nonparental time are complements, it is optimal to raise them in tandem if g falls. A parent will rely more heavily on nonparental time if the two become more substitutable and if the opportunity cost of parental time is high (e.g., $w > p$). Hence, $\psi > 0$ yields relatively modest increases in m, which require in turn smaller declines in n. In fact, as ψ approaches one, the parent purchases enough new nonparental care to support a higher n after g falls, e.g., $\mathrm{d}n/\mathrm{d}g < 0$.

It is particularly instructive to consider the case where $\mathrm{d}n/\mathrm{d}g = 0$. As shown in Online Appendix E, ψ is then bounded below such that $\psi > (1 + \xi)^{-1}$. Estimates from Blau and Currie (2006) indicate that, during the workday (prior to the pandemic), school-age children spent 1.3 hours in (private) nonparental care per hour with a parent. Setting $\xi = 1.3$ yields a *lower bound* of $\psi > 0.435$. More generally, Online Appendix E shows that for a sufficiently small but strictly

[34] In evaluating how $\mathrm{d}n/dg$ varies with ψ, we adjust μ to ensure that the initial choice of x is the same.

positive $\mathrm{d}n/\mathrm{d}g$, the bound $\psi > (1+\xi)^{-1}$ will obtain. Our estimates do suggest that $\mathrm{d}n/\mathrm{d}g$ is quite low. The return of in-person instruction represents an average of 33 weekly hours of on-site time (U.S. Department of Education, 2008), but hours worked rose by just around two (see Table 3). Therefore, estimates from Section 3 imply $\mathrm{d}n/\mathrm{d}g = 2/32 \approx 0.06$. In summary, the pandemic-era data, as seen through this model, point to significant substitutability between forms of care.

While this exercise serves to highlight the broader implications of our empirical results, a few qualifications should be noted. First, the bound on ψ is derived in the absence of other means of adjusting to shifts in g, most notably telework. A more extensive quantitative analysis of a model with nonparental care and telework is worthwhile but left for future research. Second, whereas remote instruction posed unique demands in 2020-21, time allocated to childcare in "normal" times is more diffused across academic supervision, extracurricular activities, and other tasks, some of which may require more parental inputs (see Ramey and Ramey, 2010).[35] Therefore, a more general model of the development production function would not treat adults' time as a homogeneous bloc. Notwithstanding these points, we see estimates in Berlinski et al. (2023) as broadly supporting the robustness of our conclusions. This paper studies a sample of *pre*school children—a population for whom parental time is thought to be particularly crucial to development—and still finds $\psi = 0.9$ given a production function like that used above.[36]

We conclude by highlighting that the shape of the development production function (e.g., ψ) will influence the labor market response to certain policies and shocks. First, the substitutability between forms of care mediates the impact of adolescent development policies. The federal government's role in adolescent care grew substantially with the pandemic-era expansion of the Child Care and Development Fund (U.S. Dept. of Health and Human Services, 2021). The \$40 billion of subsidies to families of school-age children effectively lowered the price of nonparental care. As one might anticipate, the demand for nonparental care rises more at higher ψs, e.g., if nonparental time is more substitutable for parental time. Therefore, the impact of the program on childcare and parental labor supply again hinges on ψ (see Online Appendix E).[37]

Second, the degree of substitution between forms of care may shape cyclical labor market dynamics. Consider for instance an increase in aggregate productivity that leads to higher wage offers. Insofar as this is a temporary shock, the hours response reflects the Frisch elasticity. The latter is, in turn, mediated by ψ: parents will substitute away from childcare toward market work

[35] See Hurst (2010) and Sacks and Stevenson (2010) for assessments of Ramey and Ramey's interpretation. For another related and insightful discussion of parental time use, see Guryan, Hurst, and Kearney (2008).
[36] We are not aware of comparable estimates for school-age children.
[37] Other, more targeted interventions, including mentoring and counseling, can be thought of as elevating the productivity of nonparental care (see Cunha et al. (2006) for a review). This policy could be captured by introducing m- and x-augmenting technologies into equation (17) and considering a perturbation to the latter.

to a greater extent if nonparental time is a close substitute for parental time. Again, see Online Appendix E for details.

6. Conclusion

This paper has presented new evidence on the response of parental labor supply, and time use more generally, to the closure of schools to on-site instruction. With a full suite of controls for unobserved heterogeneity, we do not detect any labor supply response. Even if we omit these controls, though, the labor supply responses represent a small fraction of the more than 30 hours of childcare time "lost" with the suspension of in-person instruction. Across different samples, and given a more streamlined specification, a shift from fully virtual to in-person generally implies an increase in hours worked of around two per week and rarely more than four. The paper uses a simple model of parental time allocation to formalize why these responses are unexpectedly small. Extensions to incorporate telework and nonparental care can help bridge the gap between the model and our regression estimates. We provide some evidence in the ATUS that working from home while supervising children or relying on non-parental private care may be potentially important means of coping with school closures.

Our exploration of the roles of telework and nonparental care is limited by the small sample sizes in the ATUS, and by the paucity of direct measurement of parents' utilization of non-parental care.[38] We hope our work stimulates further efforts to measure these activities. While the pandemic era was unprecedented, we suspect it highlighted the importance of margins that operate more generally. For example, adjustments to nonparental care can enable households to respond to changes in the demand for market hours. Relatedly, shifts in the composition of the household—for instance, a grandparent or an older child who moves in—can have implications for parental labor supply.

[38] One notable exception, Calarco et al. (2021) provide a recent effort to measure non-parental private childcare – either hired help or help from extended relatives – during the pandemic.

References

Adda, Jérôme, Christian Dustmann, and Katrien Stevens. 2017. “The Career Costs of Children.” *Journal of Political Economy*, 125(2): 293–337.

Alon, Titan, Matthias Doepke, Jane Olmstead-Rumsey, and Michele Tertilt. 2020. “This Time It’s Different: The Role of Women’s Employment in a Pandemic Recession.” NBER Working Paper No. 27660.

Atalay, Enghin. 2023. “Time Use Before, During, and After the Pandemic.” Mimeo, Federal Reserve Bank of Philadelphia.

Barkowski, Scott, Joanne Song McLaughlin, and Yinlin Dai. 2021. “Young Children and Parents’ Labor Supply During COVID-19.” Mimeo, University of Buffalo.

Berlinski, Samuel, Maria Marta Ferreyra, Luca Flabbi and Juan David Martin. 2023. “Child Care Markets, Parental Labor Supply and Child Development.” *Journal of Political Economy*, forthcoming.

Berthelon, Matias, Diana Kruger, and Melanie Oyarzún. 2023. “School Schedules and Mothers’ Employment: Evidence from an Education Reform.” *Review of Economics of the Household,* 21: 131-171.

Black, Dan A., Natalia Kolesnikova, and Lowell J. Taylor. 2014. “Why Do so Few Women Work in New York (and So Many in Minneapolis)? Labor Supply of Married Women Across US Cities.” *Journal of Urban Economics*, 79: 59–71.

Blau, David and Janet Currie. 2006. “Pre-school, Day Care, and After-School Care: Who’s Minding the Kids?” in Eric A. Hanushek and Finis Welch (eds.), *Handbook of the Economics of Education Volume 2,* Elsevier.

Calarco, Jessica, Max Coleman, and Andrew Halpern-Manners. 2021. “Mechanisms of Stratification in In-Person Instruction in the Wake of COVID-19.” SocArXiv Papers.

Cascio, Elizabeth U. 2009. “Maternal Labor Supply and the Introduction of Kindergarten into American Public Schools.” *Journal of Human Resources*, 44(1): 140–170.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber. 2011. “Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins.” *American Economic Review Papers & Proceedings*, 101(3): 471–475.

Cohen, Rachel. 2020. “Some Teachers Are Being Required to Come to School—to Teach Virtually,” The Intercept. Retrieved December 2022 from https://theintercept.com/2020/08/28/coronavirus-schools-teachers-remote/.

Collins, Caitlyn, Leah Ruppanner, Liana Christin Landivar, and William J. Scarborough. 2021. “The Gendered Consequences of a Weak Infrastructure of Care: School Reopening Plans and Parents’ Employment During the COVID-19 Pandemic.” *Gender and Society*, 35(2), 180–193.

Contreras, Dane, and Paulina Sepúlveda. 2017. "Effect of Lengthening the School Day on Mother's Labor Supply." *The World Bank Economic Review*, 31(3): 747–776.

Cowan, Benjamin. 2023. "Time Use, College Attainment, and The Working-from-Home Revolution." NBER Working Paper No. 31439.

Cunha, Flavio, James J. Heckman, and Susanne M. Schennach. 2010. "Estimating the Technology of Cognitive and Non-Cognitive Skill Formation." *Econometrica*, 78(3): 883–931.

Cunha, Flavio, James J. Heckman, Lance Lochner, and Dimitriy Masterov. 2006. "Interpreting the Evidence on Life Cycle Skill Formation," in Eric Hanushek and Finis Welch, eds., *Handbook of Economics of Education*. Amsterdam: North Holland.

Dee, Thomas, Elizabeth Huffaker, Cheryl Phillips, and Eric Sagara. 2022. "The Revealed Preferences for School Reopening: Evidence from Public-School Disenrollment." *American Educational Research Journal*, forthcoming.

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

Dingel, Jonathan and Brent Neiman. 2020. "How Many Jobs Can be Done at Home?" *Journal of Publc Economics*: 189: 104235.

Dong, Ensheng, Hongru Du, and Lauren Gardner. 2020. "An Interactive Web-Based Dashboard to Track COVID-19 in Real Time," *The Lancet Infectious Diseases*, 20(5): 533–534.

Education Week. 2020. "School Districts' Reopening Plans: A Snapshot" (July 15). Retrieved December 2022 from https://www.edweek.org/ew/section/multimedia/school-districts-reopening-plans-a-snapshot.html.

Elder, Laurel and Steven Greene. 2013. *The Politics of Parenthood: Causes and Consequences of the Politicization and Polarization of the American Family*. State University of New York Press: Albany.

Elder, Laurel and Steven Greene. 2021. "A Recipe for Madness: Parenthood in the Era of Covid-19." *Social Science Quarterly*, 102(5): 2296–2311.

Felfe, Christina, Michael Lechner, and Petra Thiemann. 2016. "After-school Care and Parents' Labor Supply," *Labour Economics*, 42: 64–75.

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren and Michael Westberry. Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS, 2022. https://doi.org/10.18128/D030.V10.0.

Frazis, Harley and Jay C. Stewart. 2004. "What Can Time-use Data Tell Us about Hours of Work?" *Monthly Labor Review*, December: 3–9.

Frazis, Harley and Jay C. Stewart. 2014. "Is the Workweek Really Overestimated?" *Monthly Labor Review*, June.

Furman, Jason, Melissa Schettini Kearney, and Wilson Powell. 2021. “The Role of Childcare Challenges in the US Jobs Market Recovery During the COVID-19 Pandemic.” National Bureau of Economic Research Working Paper No. 28934.

Garcia, Kairon Shayne and Benjamin Cowan. 2022. “The Impact of U.S. School Closures on Labor Market Outcomes during the COVID-19 Pandemic,” NBER Working Paper No. 29641.

Gelbach, Jonah B. 2002. “Public Schooling for Young Children and Maternal Labor Supply.” *American Economic Review*, 92(1): 307–322.

Goldin, Claudia. 2022. “Understanding the Economic Impact of COVID-19 on Women.” *Brookings Papers on Economic Activity*, Spring: 65–110.

Görlitz, Katja, and Marcus Tamm. 2020. “Parenthood, Risk Attitudes and Risky Behavior.” *Journal of Economic Psychology*, 79: 1–20.

Grossmann, Matt, Sarah Reckhow, Katharine Strunk, and Meg Turner. 2021. “All States Close but Red Districts Reopen: The Politics of In-Person Schooling during the COVID-19 Pandemic.” *Educational Researcher*, 50(9): 637–648.

Guner, Nezih, Remzi Kaygusuz, and Gustavo Ventura. 2020. “Child-Related Transfers, Household Labor Supply and Welfare.” *Review of Economic Studies*, 87(5): 2290–2321.

Guryan, Jonathan, Erik Hurst, and Melissa Kearney. 2008. “Parental Education and Parental Time with Children.” *Journal of Economic Perspectives*, 22(3): 23–46.

Hansen, Benjamin, Joseph Sabia, and Jessamyn Schaller. 2022. “Schools, Job Flexibility, and Married Women's Labor Supply: Evidence From the COVID-19 Pandemic.” National Bureau of Economic Research Working Paper 29660.

Hartney, Michael and Leslie Finger. 2021. “Politics, Markets, and Pandemics: Public Education’s Response to COVID-19.” *Perspectives on Politics*, 20(2): 457–473.

Heggeness, Misty L. 2020. “Estimating the Immediate Impact of the COVID019 Shock on Parental Attachment to the Labor Market and Double Bind of Mothers,” *Review of Economics of the Household*, 18(4): 1053-1078.

Heggeness, Misty L. and Palak Suri. 2021. “Telework, Childcare, and Mothers’ Labor Supply.” Opportunity and Inclusive Growth Institute Working Paper No. 52.

Hofferth, Sandra L., Flood, Sarah M., Sobek, Matthew, and Backman, Daniel. 2020. American Time Use Survey Data Extract Builder: Version 2.8 [dataset] College Park, MD: University of Maryland and Minneapolis, MN: IPUMS, 2020. https://doi.org/10.18128/D060.V2.8.

Hurst, Erik. 2010. “Comment on: The Rug Rat Race.” *Brookings Papers on Economic Activity*, Spring: 177–184.

Kerry, Nicholas, et al. 2022. “Experimental and Cross-cultural Evidence that Parenthood and Parental Care Motives Increase Social Conservatism.” *Proceedings of the Royal Society B*, 289(1982).

Kleven, Henrik, Camille Landais, and Jakob Egholt Søgaard. 2019. "Children and Gender Inequality: Evidence from Denmark." *American Economic Journal: Applied Economics*, 11: 181–209.

Kurmann, Andre and Etienne Lalé. 2023. "School Closures and Effective In-Person Learning during COVID-19." *Economics of Education Review*, 95: 102422.

Jung, Carrie. 2020. New State Guidance: Remote Teaching Should Happen from the Classroom. wbur.org, 21 August, https://www.wbur.org/news/2020/08/21/massachusetts-remote-learning-teachers-in-classrooms-dese-guidance.

Lofton, Olivia, Nicolas Petrosky-Nadeau, and Lily Seitelman. 2021. "Parents in a Pandemic Labor Market," Federal Reserve Bank of San Francisco Working Paper 2021-04.

Lozano-Rojas, Felipe, Xuan Jiang, Laura Montenovo, Kosali I. Simon, Bruce A. Weinberg, and Coady Wing. 2020. "Is the Cure Worse Than the Problem Itself? Immediate Labor Market Effects of COVID-19 Case Rates and School Closures in the U.S." NBER Working Paper No. 27127.

Marianno, Bradley, Annie Hemphill, Ana Paula Loures-Elias, Libna Garcia, Deanna Cooper, and Emily Coombes. 2022. "Power in a Pandemic: Teachers' Unions and Their Responses to School Reopening." *AERA Open*: 8.

Mumford, Karen, Antonia Parera-Nicolau, and Yolanda Pena-Boquete. 2020. "Labour Supply and Childcare: Allowing Both Parents to Choose." *Oxford Bulletin of Economics and Statistics,* 82(3): 577-602.

Musaddiq, Tareena, Kevin Stange, Andrew Bacher-Hicks, and Joshua Goodman. 2022. "The Pandemic's Effect on Demand for Public Schools, Homeschooling, and Private Schools." *Journal of Public Economics*, 212: 104710.

Pabilonia, Sabrina Wulff, and Victoria Vernon. 2023. "Who is Doing the Chores and Childcare in Dual-Earner Couples During the COVID-19 Era of Working from Home?" *Review of Economics of the Household*, 21 519–565.

Padilla-Romo, Maria and Francisco Cabrera-Hernández. 2019. "Easing the Constraints of Motherhood: The Effects of All-Day Schools on Mothers' Labor Supply." *Economic Inquiry*, 57: 890–909.

Parolin, Zachary and Emma Lee. 2021. "Large Socio-economic, Geographic and Demographic Disparities Exist in Exposure to School Closures." *Nature Human Behaviour*, 5: 522–528.

Planty, Michael, William Hussar, Thomas Snyder, Stephen Provasnik, Grace Kena, Rachel Dinkes, Angelina KewalRamani, and Jana Kemp. 2008. "The Condition of Education, 2008. NCES 2008-031." *National Center for Education Statistics*.

Ramey, Garey and Valerie Ramey. 2010. "The Rug Rat Race." *Brookings Papers on Economic Activity*, Spring: 129–176.

Sacks, Daniel and Betsy Stevenson. 2010. "Comment on: The Rug Rat Race." *Brookings Papers on Economic Activity*, Spring: 184–196.

Solon, Gary, Steven J. Haider, and Jeffrey M. Wooldridge. 2015. "What Are We Weighting For?" *Journal of Human Resources,* 50(2): 301–316.

U.S. Department of Education. 2008. "Average number of hours in the school day and average number of days in the school year for public schools, by state: 2007–08," `https://nces.ed.gov/surveys/sass/tables/sass0708_035_s1s.asp`.

U.S. Department of Health and Human Services. 2021. Information Memorandum: American Rescue Plan Act Child Care Stabilization Funds, Office of Child Care.

Wu, Pinghui, Vincent Fusaro, and H. Luke Shaefer. 2022. "Government Transfers and Consumer Spending among Households with Children during COVID-19." Federal Reserve Bank of Boston Working Paper No. 22-17.

Yang, Yining Milly, Emma Zang, and Jessica McCrory Calarco. 2022. "Patterns in Receiving Informal Help with Childcare Among U.S. Parents During the COVID-19 Pandemic." SocArXiv Papers.

Online Appendix to "School Closures, Parental Labor Supply, and Time Use"

Enghin Atalay, Ryan Kobler, and Ryan Michaels

A. The spread of the pandemic and instruction format

This appendix reviews evidence on the correlation between COVID-19 transmission and the choice of instruction format. Since the area's in-person share is the object of interest, the analysis is carried out at the area (by month) level. The full set of local areas is used. Throughout, the sample period is the 2020-21 school year. The latter choice eliminates the first few months of the pandemic when in-person instruction was suspended in anticipation of the spread of the disease and was, in fact, approximately independent of observed cases and fatalities. Therefore, the correlation between COVID-19 transmission and school policy is likely to be highest in the shorter subsample we consider here.[1]

Our focus here concerns, specifically, the potential response of school policy to the evolution of the pandemic. To that end, we run regressions with area fixed effects to isolate *within-area* co-movement of COVID-19 transmission and the in-person share. To be sure, the influence of the pandemic on the path of school policy may be mediated by (fixed) local attributes. Therefore, we will also consider a specification that involves a full slate of interactions between certain attributes and indicators of the spread of COVID-19. Appendix B presents a fuller exploration of the *cross-area* variation in average school policy.

Table A1 reports our estimates. The outcome is Parolin and Lee's in-person share. The initial batch of regressors consists only of monthly COVID-19 cases and fatalities as well as area and month fixed effects. We highlight three results. First, the *within* R-squared of less than 0.04 indicates that these regressors account for a small portion of overall variation in in-person shares. Second, a higher number of cases and fatalities both imply lower in-person shares. Specifically, one more case per 100 area residents is associated with a three percentage point reduction in the in-person share, and one higher fatality per 10,000 residents implies a nearly one percentage point reduction. These estimates are statistically significant but modest in size. To illustrate, a two standard deviation movement in cases implies less than one-fifth of a standard deviation movement in in-person shares. A comparable shift in fatalities has a still smaller impact.

Next, we interact cases and fatalities with a vector of local attributes. The latter consists, first, of various demographic controls. A large body of research finds that the pandemic led to sharp increases in mortality among certain groups, especially Hispanics, non-White individuals, and the noncollege educated (see, i.e., Case and Deaton, 2021). There may be a greater demand for social distancing, including virtual instruction, in areas where such groups are highly represented. Accordingly, the vector of attributes includes each group's share in the local population. Mortality rates may have also risen more in areas of high density.[2] To capture the latter, our list of attributes includes the share of an area's population in a city center.

Our set of local attributes also consists of policy-relevant institutions and political attitudes. For example, areas with more unionized education sectors returned to in-person instruction later in the 2020-21 school year. In addition, areas that favored the Republican party generally chose more on-site instruction. We use two indicators of Republican party strength: Donald Trump's share of the area's 2016 presidential vote and the presence of a Republican governor as of January 2020. With the latter two as exceptions, all

[1] The results are quite similar if we include fall 2021, too.

[2] See, for instance, Almagro and Orane-Hutchinson (2020). Carozzi et al. (2021) have challenged this claim, though.

attributes are computed with CPS data pooled over the pre-pandemic period, 2015-19. Donald Trump's 2016 vote share is from the MIT Election Data and Science Lab. The partisan identity of governors in 2020 is from ballotpedia.org.[3]

The second column in Table A1 reports the coefficient estimates for this expanded specification. To aid interpretation, each attribute is expressed as a deviation from the national average. Thus, the coefficient estimates on cases and deaths (first two rows) indicate the response of the in-person share when all attributes are evaluated at their mean. These estimates are negative, statistically significant, and similar to those in the first column. A positive (negative) coefficient on an interaction term implies that the associated attribute mitigates (amplifies) the decline in in-person instruction that coincides, on average, with higher cases and deaths.

A few results emerge from the table. To start, the introduction of the interaction terms elevates the within R-squared, but the latter remains around just 0.12. Our takeaway from this result is that month to month variation in the state of the pandemic had a limited impact on the evolution of instruction format.

To the extent that the spread of COVID-19 shaped school policy, it did so in a heterogeneous manner. The political identity of the local area played a notable role. For instance, in an area where Trump's share was one standard deviation (or, 14.4 percentage points) below the mean, one more confirmed case per 100 residents was associated with a decrease in the in-person share of $-0.036 - 0.144 \times 0.145 = 5.7$ percentage points. By contrast, the in-person share falls just 1.5 percentage points in an area with one-standard-deviation more Trump support. The same contrast applies to higher COVID-19-related fatalities.

Other notable local attributes include the nonwhite and Hispanic shares, but their impact is less straightforward. On the one hand, in areas with high non-White and Hispanic shares, the in-person share declines *by less* when COVID-19 cases rise. Indeed, a one standard deviation shift in either of the latter has nearly the same impact as a comparably scaled shift in the Trump share. On the other hand, when fatalities rise, the in-person share declines *by more* in areas with high minority shares.

Finally, the remaining attributes play a less consistently significant role. The presence of a Republican (GOP) governor echoes the effect of the Trump share when interacted with COVID-19 cases but not fatalities. By the same token, the unionized share of teachers and the city-center share of population enter significantly in some interactions but not others. The noncollege share is insignificant throughout.

In sum, school policy in more heavily Democratic areas was more sensitive to the trajectory of the pandemic. However, the overall explanatory power of COVID-19 case and death counts is rather modest.

[3] On the 2016 vote share, see the "U.S. President 1976-2020" data file at `https://electionlab.mit.edu/data`. The list of present-day governors is at `https://ballotpedia.org/Governor_(state_executive_office)`. The latter was accessed December 2022. For a handful of cases where the governorship had turned over since 2020, we separately looked up the party of the governor in 2020.

Table A1: The Spread of COVID-19 and School Policy

Coefficient		In-person share (SafeGraph)	
Monthly cases / 100		-0.030***	-0.036***
		[0.005]	[0.005]
Monthly deaths / 10,000		-0.009***	-0.009***
		[0.002]	[0.003]
Monthly cases ×	Trump share		0.144***
			[0.040]
	GOP governor		0.042***
			[0.009]
	Teacher union		-0.018
			[0.018]
	Hispanic		0.144***
			[0.023]
	Nonwhite		0.157***
			[0.040]
	Noncollege		-0.050
			[0.044]
	City center		0.040**
			[0.017]
Monthly deaths ×	Trump share		0.056***
			[0.021]
	GOP governor		-0.007*
			[0.004]
	Teacher union		-0.020**
			[0.008]
	Hispanic		-0.024*
			[0.012]
	Nonwhite		-0.042**
			[0.017]
	Noncollege		0.009
			[0.024]
	City center		-0.000
			[0.010]
Number of obs.		4,293	4,293
Within R^2		0.036	0.123

Note: The sample period is the 2020-21 school year, and the unit of analysis is the local area × month. Area fixed effects are included. Other than COVID-19 cases and deaths, regressors are expressed as deviations from the U.S. average. See text for more. Standard errors are clustered at the geographic area level. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

B. Placebo test results

Table B1 reports the results of the placebo test described near the end of Section 3.1. We relate *pre*-pandemic labor supply to the average share of the school year for which schools were open in the pandemic. We present results for all three of our measures of school formats: SafeGraph, CSDH, and Burbio. The latter two are available only for the 2020-21 school year (which spans September 2020-May 2021). Therefore, for the sake of comparability, we report SafeGraph-based estimates that average in-person instruction shares over just the 2020-21 school year as well as over all of 2020-21.

Table B1 lists the results from this exercise. Since we discussed the SafeGraph-based results in the main text, our comments here pertain mainly to our other two data sources. Consider first the results for women based on CSDH data. To interpret these, recall that the idea behind the regression is to ask if differences across areas in 2020-21 in-person shares predict differences in mothers' relative hours in 2015-19 (that is, relative to the hours of childless women in her area at that time). Thus, we read Table B1 to say that, in the five years prior to the pandemic, mothers' relative labor supply in an area with full-time in-person instruction was nearly 1.5 weekly hours higher than in an area with full-time virtual instruction. This result is somewhat smaller than its SafeGraph-based counterpart over the school year. Estimates from Burbio show a weaker, but still statistically significant, relationship between pandemic-era policy and pre-pandemic labor supply outcomes.

Next, we turn to the results for men. The estimates are uniformly smaller than the labor supply responses uncovered in the pandemic period (see Table 3), although one often cannot reject equivalence. Thus, as we saw for women, the placebo estimates are in the same neighborhood as those obtained over the 2020-21 sample. The placebo estimates are less precisely estimated, though. In particular, the SafeGraph-based estimates are only marginally significant. However, fathers' labor supply response based on CSDH data is somewhat larger, and more precisely estimated, than what we observed in the pandemic period. By contrast, estimates off Burbio data are insignificant in both periods.

Table B1: Placebo Test

	SafeGraph		CSDH	Burbio
Coefficient		Women		
In-person × kids, ψ	3.353***	2.007***	1.427***	1.132**
	[0.877]	[0.558]	[0. 444]	[0.442]
Period of policy	All 20-21	School 20-21	School 20-21	School 20-21
Number of obs.	1,351,083	1,351,083	1,254,179	1,245,826
Coefficient		Men		
In-person × kids, ψ	1.386*	0.855*	0.737*	0.406
	[0.851]	[0.525]	[0.450]	[0.418]
Period of policy	All 20-21	School 20-21	School 20-21	School 20-21
Number of obs.	1,284,357	1,284,357	1,191,245	1,184,507

Note: This table estimates a version of equation (1) on CPS data 2015-19. Relative to equation (1), the policy variable is the pandemic-era mean. The "period of policy" refers to the specific years over which the mean is taken: "All 20-21" includes calendar years 2020 and 2021 (exclusive of June-August), whereas "School 20-21" covers only September 2020 – May 2021. Standard errors are clustered at the geographic area level. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

Overall, these estimates suggest that in-person shares capture more general, and fundamental, forces behind parental labor supply. To this extent, we also expect in-person shares to be related to other (pre-pandemic) outcomes that are likely co-determined with parental hours and, therefore, shaped by these same forces. The main text highlights two such outcomes: childcare prices and commute-to-work times.

The connection between these outcomes and pandemic-era school policy runs, in part, through their association with local partisan affiliation. Figure B1 shows that lower commute times and childcare prices are each associated with higher support for Donald Trump. (We report our data sources momentarily.) This result is significant because the vote share for President Trump is strongly and positively correlated with in-person instruction shares. It follows that lower commute times and childcare prices are likely to predict higher in-person shares. Figure B2 confirms this claim.

Table B2 offers a simple statistical summary of these points. We regress a local area's average in-person share over the period 2020-21 on three variables: Trump's share of the 2016 presidential vote, the average commute time in 2015-19, and average pre-pandemic childcare prices. Donald Trump's share is from the MIT Election Data and Science Lab (see Appendix A). Mean commute times are taken from the Census Bureau's county-level tabulations of the 2015-19 American Community Survey 5-year estimates (Manson et al., 2022). Childcare prices are compiled by the Women's Bureau of the U.S. Department of Labor. The regression sample consists of 404 local areas for which center-based and (smaller-scale) family-based prices are available from the Women's Bureau data. Finally, the regressions are fit to childcare prices averaged across both types of care and commute times averaged over men and women. The figures, which report separate scatter plots by type of care and gender, indicate that little detail is lost if the data are pooled.[4]

The estimates in Table B2 corroborate, and extend, the evidence in Figures B1 and B2. First, childcare prices and commute times are each negatively, and significantly, associated with in-person shares. A price increase of $100/week implies a nearly 20 percentage point lower in-person share. In addition, a 10 minute longer commute is associated with a 10 percentage-point lower in-person share. However, Trump's support is the most significant predictor of in-person shares among the three regressors. Indeed, when Trump's vote share is added to the regression, the coefficients on commute time and childcare prices are halved. Thus, the connection between these two factors and school policy is mediated, in part, by the factors' association with partisan affiliation. Nevertheless, the other two factors do account for some portion of the variance in in-person shares conditional on Trump's support.

Table B2: Correlates of In-person Instruction

Coefficient	SafeGraph in-person share					
Trump vote share	0.696***			0.581***	0.652***	0.573***
	[0.022]			[0.028]	[0.022]	[0.025]
Childcare price / 100		-0.194***		-0.082***		-0.064***
		[0.014]		[0.011]		[0.010]
Commute time / 10			-0.102***		-0.056***	-0.042***
			[0.012]		[0.007]	[0.007]
Number of obs.	404	404	404	404	404	404
R^2	0.662	0.374	0.161	0.712	0.708	0.735

Note: See text for description of data sources. Note that the childcare price is expressed in hundreds of dollars, and commute time is expressed in tens of minutes. Standard errors are robust to heteroskedasticity. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

[4] We drop counties if their pre-pandemic childcare prices were imputed based on state-level prices.

Figure B1: Childcare Prices, Commute Times, and Partisan Affiliation

Panel A: Childcare Prices

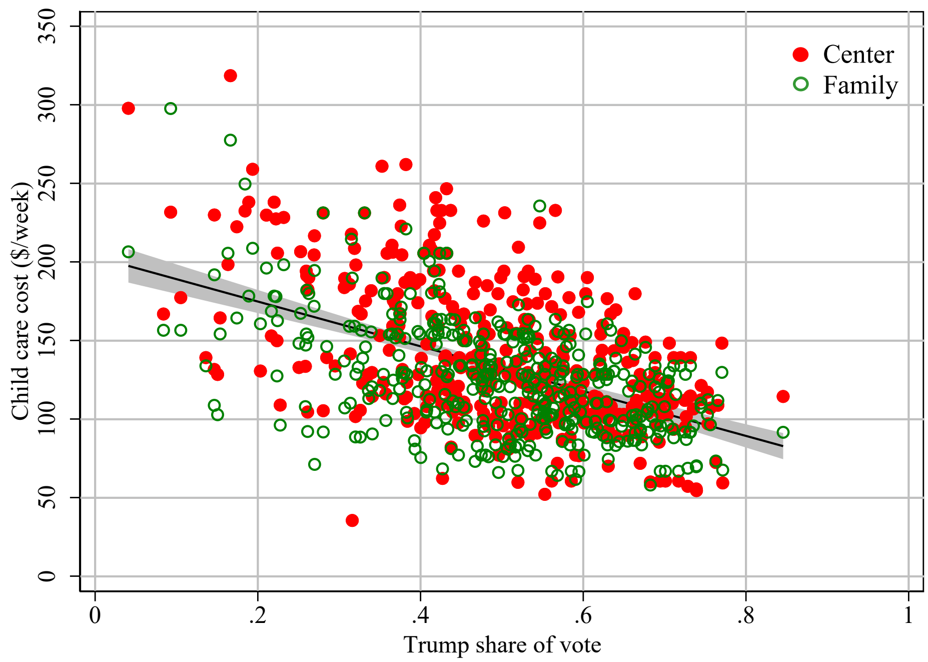

Panel B: Commute Times

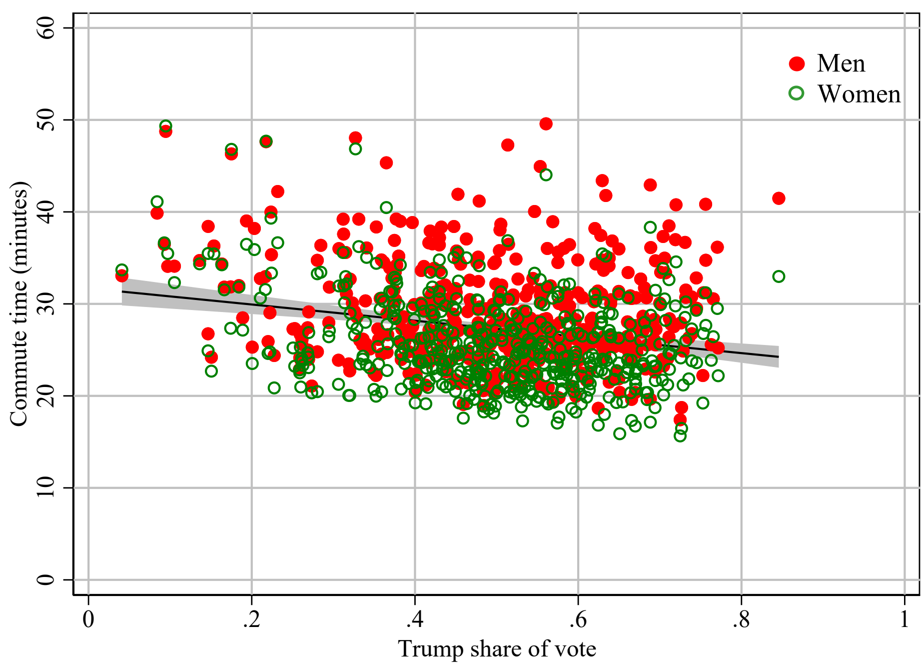

Note: The line of best fit in each panel reflects a regression on the pooled sample, e.g., in Panel A, separate scatter plots are shown by type of childcare center, but the regression line is fit to the average price across both types. Similarly, the regression line in Panel B is fit to the average commute time across men and women. The gray shaded region shows the 95% confidence band. See text for a description of the data sources.

Figure B2: Childcare Prices, Commute Times, and In-person Shares

Panel A: Childcare Prices

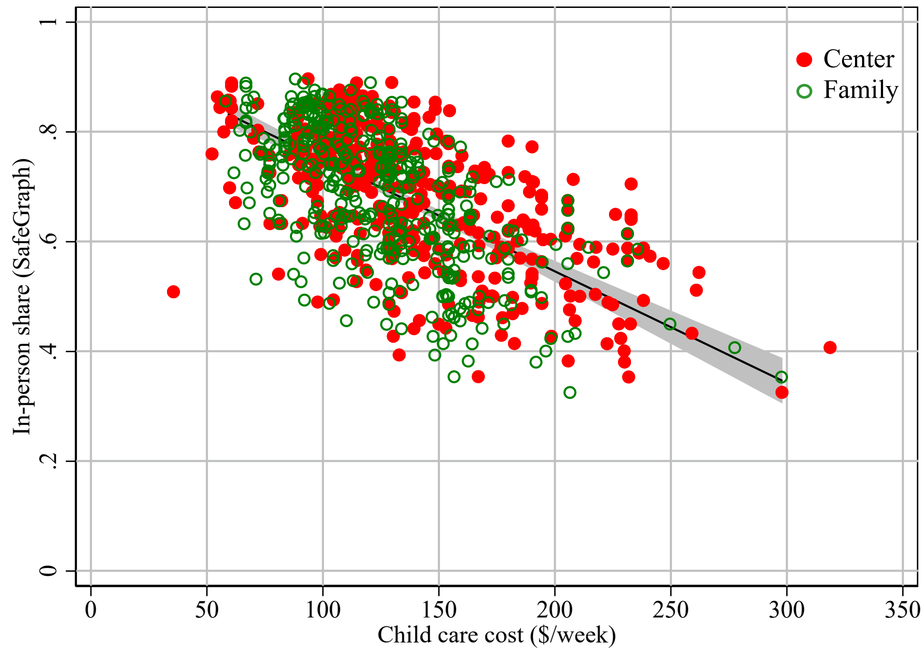

Panel B: Commute Times

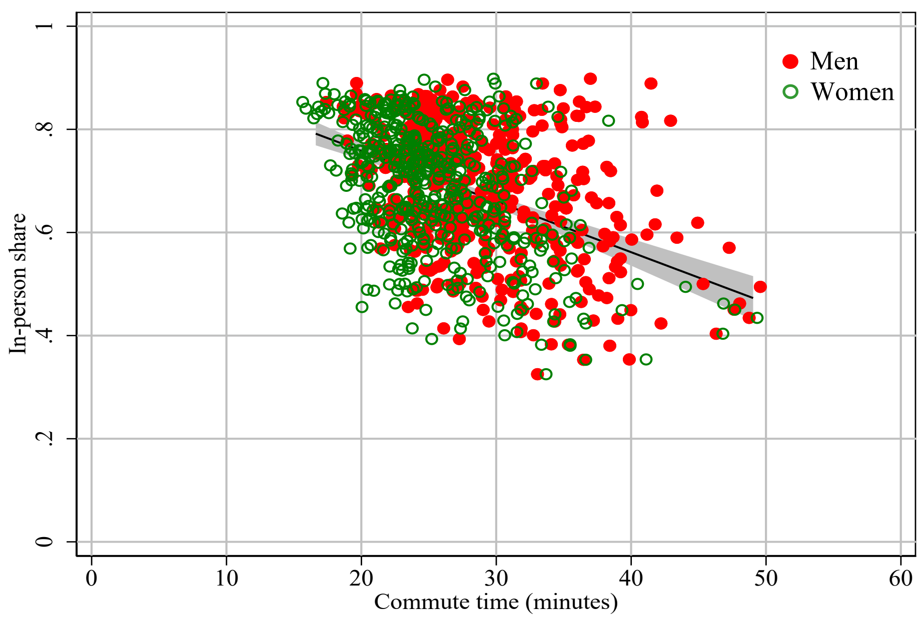

Note: See notes to Figure B1.

To follow up on these results, we also examined if the connection between childcare prices and the Trump share reflects a partisan role in childcare regulations. This analysis is done at the *state* level since childcare is regulated by state authorities. For the purpose of regression analysis, we zero in on one regulation: the maximum child-staff ratio (see also Kimmel, 1998). A higher maximum enables a childcare center to operate, in principle, with fewer staff and, therefore, at lower cost. We draw on a database of state regulations maintained by the National Center on Early Childhood Quality Assurance and measure the maximum child-staff ratio as of 2017, the midpoint of our pre-pandemic sample. Across the states, this maximum varies from 10 to 25 with an interquartile range of 15 to 20.[5]

We find a statistically significant, but quantitatively modest, connection between partisanship, childcare regulations, and childcare prices. An increase in Trump's share of 10 percentage points implies one more child per staff member is permitted under state law. This latter is statistically significant but represents a fairly small share of the variation in child-staff ratios. By the same token, the allowance for one more child per staff member has a limited impact on weekly childcare prices, which fall by 2-3 percent (the higher of the two pertains to center-based rather than family-based care).

As a final exercise, we document how commute times and childcare prices are related to pre-pandemic labor supply. Table B3 reports the correlations by gender, conditional on the same set of demographic covariates used in Table B1. These regressions are run on individual-level CPS data in 2015-19 across the 404 local areas for which we have childcare prices.[6] The results differ markedly across men and women. Table B3 shows that commute time and childcare prices are each individually significant correlates of *maternal* hours worked. Specifically, weekly hours decline by around one per \$100 increase in the weekly childcare price and per 10-minute increase in commute time. However, we find that neither commute times nor childcare prices are relevant to *paternal* hours worked.

Table B3: Childcare Prices, Commute Times, and Hours Worked

	Weekly hours worked					
Coefficient	Women			Men		
Childcare cost	-1.129***		-0.915***	-0.354		-0.361
× kids	[0.331]		[0.300]	[0.315]		[0.364]
Commute time		-0.849***	-0.577*		-0.076	0.019
× kids		[0.237]	[0.264]		[0.239]	[0.291]
Number of obs.	1,069,053	1,069,053	1,069,053	1,015,075	1,015,075	1,015,075

Note: The childcare price is the mean of center- and family-based prices and expressed in hundreds of dollars. Commute time is the mean among employed adults ages 21-59 and expressed in tens of minutes. Standard errors are clustered at the geographic area level. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

B. Sensitivity analysis and additional results from the CPS

This appendix reports additional results on the labor supply response to variation in the in-person share. First, we estimate the role of the extensive margin in total hours adjustment. Next, we report results from a battery of sensitivity tests. Specifically, we examine the implications of the age range of adult respondents; the geographic coverage of the sample; the use of industry and occupation controls; and

[5] These data are available from ICPSR. See `https://www.icpsr.umich.edu/web/ICPSR/studies/37700`. We use data from the 2017 licensing study.
[6] We have confirmed that the placebo test also fails in this subsample. Indeed, the estimate of ψ (see Table B1) in this subsample is *higher* among fathers (and hardly affected among mothers).

alternative measures of the in-person share. As a way to synthesize some of these results, we then trace through in detail how our estimates relate to those in Garcia and Cowan (2022), who use a wider age range, a narrower geographic coverage, and industry and occupation controls. As a final exercise, we rerun our main specification but take the *household* as the unit of analysis. This estimation yields the response of the household's total hours of work to variation in instruction format.

C.1 Extensive margin

Tables C1 and C2 report results for employment (rather than weekly hours). Table C1 is the employment analogue to Table 3 in the main text, which presents estimates of equations (1)-(2) based on SafeGraph data. Table C2 is the employment analogue to Table 4 in the main text, which is based on our two alternative measures of in-person shares from Burbio and the COVID-19 School Data Hub (CSDH). (We return to discuss Burbio and CSDH data at length later in this appendix.)

Estimates in Table C1 mirror the results for hours in the main text. First, conditional on controls for unobserved heterogeneity, the parental labor supply response is stable across sample periods. Second, the extensive margin accounts for the labor supply response of mothers but not for fathers. Given a 37-hour week among employed mothers, the 5.6 percentage-point gain in the maternal employment rate (see "All 20-21") implies an increase in *total* weekly hours of 2.1—a near replica of the estimate in Table 3. Third, the introduction of area-by-parental-status effects eliminates the statistical significance of these estimates.

Table C1: Employment Responses Based on SafeGraph In-person Shares

	All 20-21	School 20-21	All 20-21	School 20-21
Coefficient		Women		
In-person × kids, ψ	0.056***	0.059***	-0.005	-0.006
	[0.014]	[0.015]	[0.016]	[0.027]
Number of obs.	447,899	228,550	447,899	228,550
Coefficient		Men		
In-person × kids, ψ	0.013	0.012	0.000	-0.012
	[0.014]	[0.013]	[0.014]	[0.022]
Number of obs.	432,856	221,080	432,856	221,080
Month × parent	Yes	Yes	Yes	Yes
Area × parent	No	No	Yes	Yes

Note: "All 20-21" refers to the 2020 and 2021 calendar years save for June, July, and August. "School 20-21" refers to the period September 2020 to May 2021. Standard errors are clustered at the geographic area level. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

As we saw in regressions with hours, the response of maternal employment is estimated to be smaller if in-person shares are taken from Burbio or CSDH. Results based on the latter measures are reported in Table C2. Once more, we can confirm that the extensive margin accounts for virtually all the increase in hours reported for the analogous regression in the main text (Table 4).

Next, we present employment estimates by educational attainment and marital status in Table C3. This table is the extensive-margin counterpart to Tables 6 and 7. Here, as in our baseline regressions, the in-person share if based on SafeGraph data taken from Parolin and Lee (2021). The left panel breaks the sample by educational attainment, e.g., (four-year) college and noncollege graduates. The right panel reports regression results by marital status. In both panels, the parental employment response is estimated by the specification used in Section 3.2, which augments equation (1) with parental-status-by-month effects.

Table C2: Employment Responses Based on Alternative In-person Shares

	CSDH			Burbio		
Coefficient	Women					
In-person × kids, ψ	0.0244**	0.034***	-0.024	0.016	0.026**	0.007
	[0.012]	[0.013]	[0.020]	[0.010]	[0.011]	[0.016]
Number of obs.	211,156	211,156	211,156	211,777	211,777	211,777
Coefficient	Men					
In-person × kids, ψ	-0.001	0.001	-0.015	0.005	0.008	0.023
	[0.010]	[0.011]	[0.016]	[0.009]	[0.010]	[0.014]
Number of obs.	204,090	216,034	216,034	205,039	205,039	205,039
Month × parent	No	Yes	Yes	No	Yes	Yes
Area × parent	No	No	Yes	No	No	Yes

Note: The left panel is based on in-person shares from CSDH. The right panel is based on Burbio estimates. Throughout, the sample period is September 2020-May 2021. Standard errors are clustered at the geographic area level. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

Table C3: Employment Responses by Education and Marital Status

	Education		Marital Status		
	Noncollege	College	Married	Unmarried	Lone adult
Coefficient	Women				
In-person × kids, ψ	0.049***	0.057***	0.056***	0.059**	0.110***
	[0.018]	[0.021]	[0.018]	[0.023]	[0.026]
Number of obs.	266,258	181,641	242,743	205,156	67,592
Coefficient	Men				
In-person × kids, ψ	0.007	0.010	0.031**	-0.002	0.071
	[0.016]	[0.018]	[0.014]	[0.030]	[0.045]
Number of obs.	284,723	148,133	223,471	209,385	61,954

Note: Each column reports an estimate based on equation (1) but where the outcome is an employment indicator and month-by-parental-status effects are added to the controls. The period is calendar years 2020 and 2021 with the summer months (June-August) omitted. Standard errors are clustered at the geographic area level. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

These results echo a theme in the main text, namely, the extensive margin was active for mothers but much less so for fathers. Consider, for instance, the result for noncollege educated mothers. A shift from virtual to in-person instruction implies an increase in the county's maternal employment rate of nearly 5 percentage points. Given an average workweek of 36 hours among noncollege employed mothers, the extensive-margin response accounts for 1.76 additional weekly hours of work. This portion represents 85 percent of the estimated response of (actual) weekly hours reported in Table 6 (2.07). Other results for women in Table C3 send the same message. By contrast, among fathers, the extensive margin is inactive more often than not. The one exception is for married fathers, who experience a statistically significant increase in employment upon the return of in-person instruction. This estimate accounts for just over 70 percent of the response in (actual) weekly hours shown in Table 7.[7]

C.2 Age range of adults

Table C4 reports estimates of the parental labor supply response for two age ranges. The left panel recapitulates the estimates from our preferred sample that consists of adults ages 21-59. The first column reports the parental labor supply response, e.g., ψ, based on estimation of equation (1). In the second and third columns, additional controls for parent-specific time and area effects are introduced as in equation (2). The right panel repeats these regressions for a sample that consists of all adults ages 21 and over. Throughout, the dependent variable is weekly hours.

The table reveals three results. First, the estimated parental labor supply response based on equation (1) is at least three times larger when adults over age 59 are included in the sample. Second, the difference in estimates across samples is not nearly so large when parent-specific time effects are included in the specification. Third, when additional spatial controls are added, the parental labor supply response is indistinguishable from zero in each sample.

Table C4: Age Limit of Adults in Sample

	Ages 21-59			All ages		
Coefficient	Women					
In-person × kids, ψ	0.582*	2.360***	-0.040	1.969***	2.868***	-0.033
	[0.304]	[0.634]	[0.672]	[0.264]	[0.589]	[0.576]
Number of obs.	447,899	447,899	447,899	728,758	728,758	728,758
Coefficient	Men					
In-person × kids, ψ	0.567*	1.888***	-0.051	1.892***	2.610***	0.443
	[0.314]	[0.645]	[0.706]	[0.292]	[0.629]	[0.666]
Number of obs.	432,856	432,856	432,856	671,403	671,403	671,403
Month × parent	No	Yes	Yes	No	Yes	Yes
Area × parent	No	No	Yes	No	No	Yes

Note: The left panel reports results for our baseline sample. The right panel reports results for a sample of adults age 21 and over. Columns within each panel are differentiated by the inclusion of month-by- and/or area-by-parental status controls. Standard errors are clustered at the geographic area level. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

There is a simple interpretation of these results. With older adults in the sample, the parental labor supply response cannot be distinguished in equation (1) from the labor supply response of under-age-60 adults more generally. Put another way, the OLS estimate of ψ reflects a common component in hours shared among all adults under age 60. Table C5 illustrates this point. To equation (1), we add an interaction between p_{at} and an indicator for under age 60. If the in-person share were strictly exogenous, the latter

[7] The average workweek of employed married fathers is 43 hours. The response in the table then implies an increase in weekly hours of 1.3, whereas the total weekly hours response (Table 7) is 1.8.

interaction would be indistinguishable from zero *conditional on* the interaction between p_{at} and parental status. In fact, the under-age-60 term enters as strongly significant: a shift from fully virtual to fully in-person implies an increase in market work among *all* under-age-60 respondents of between 3-3.5 hours per week. In this context, the interaction between p_{at} and parental status now isolates the response of parental labor supply *relative to* the average response among under-age-60 respondents. This response is more akin to what we identify in the sample of adults ages 21-59. Indeed, Table C5 reveals that the coefficient on the interaction between p_{at} and parental status is comparable to what we report for the latter sample (on the left panel of Table C4).

A final point to note about Table C4 concerns the impact of the month-by-parental status effects. The introduction of these controls reduces the gap between the estimates of ψ in the ages 21-59 sample and the all-ages sample. This finding suggests that a portion of the variation identifying ψ in the all-ages sample in fact reflects aggregate trends in working-age labor supply. The latter are partially controlled for by the month-by-parental status effects.

Table C5: All Ages with Controls for Working Age Status

Coefficient	Women	Men
In-person × kids	0.770**	0.649**
	[0.323]	[0.319]
In-person × $\mathbb{I}[21 \leq \text{age} \leq 59]$	2.900***	3.547***
	[0.355]	[0.323]
Number of obs.	728,758	671,403

Note: The sample includes ages 21 and over. The "working age" refers to adults ages 21-59. Each column reports an estimation of equation (1) with two added controls: an indicator for working age, $\mathbb{I}[21 \leq \text{age} \leq 59]$; and an interaction of the latter with the in-person share. Standard errors are clustered at the geographic area level. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

C.3 Geographic coverage

Table C6 reports results for two different samples differentiated by their geographic coverage. The format of the table mirrors of that Table C4. Hence, the left panel recapitulates our baseline results for the full sample of 478 local areas. The right panel reports for the alternative sample, in this case a subsample of our baseline comprised only of the 280 counties in the CPS.

The theme of Table C6 is that this smaller subsample implies a weaker labor supply response. Consider the specification with time-by-parental status controls. The labor supply response among mothers falls from 2.4 hours per week in our preferred sample to 1.6 hours per week. Among fathers, the response is somewhat smaller and no longer statistically significant.

C.4 Industry and occupation controls

Next, we introduce controls for industry and occupation. Specifically, we include indicator variables to span 17 industries, each of which corresponds to a two-digit NAICS sector. We also include indicator variables to span 23 occupations, each of which corresponds approximately to a two-digit SOC code. In the CPS, these industry and occupation codes are available for all labor force participants.

Table C6: Geographic Coverage

	All local areas			CPS counties		
Coefficient	Women					
In-person × kids, ψ	0.582*	2.360***	-0.040	0.214	1.555*	0.127
	[0.304]	[0.634]	[0.672]	[0.390]	[0.935]	[1.019]
Number of obs.	447,899	447,899	447,899	188,204	188,204	188,204
Coefficient	Men					
In-person × kids, ψ	0.567*	1.888***	-0.051	0.129	1.089	-0.559
	[0.314]	[0.645]	[0.706]	[0.434]	[0.984]	[1.035]
Number of obs.	432,856	432,856	432,856	179,594	179,594	179,594
Month × parent	No	Yes	Yes	No	Yes	Yes
Area × parent	No	No	Yes	No	No	Yes

Note: The left panel refers to our baseline sample. The right panel is based on the sample of CPS-identified counties. Columns within each panel are differentiated by the inclusion of month-by- and/or area-by-parental status controls. Standard errors are clustered at the geographic area level. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

However, nonparticipants report industry and occupation only if they (i) are in the Outgoing Rotation Groups and (ii) have worked in the last 12 months. In practice, these restrictions severely limit the availability of industry and occupation codes, which is reported for just 3.5 percent of nonparticipants (ages 21-59). To accommodate the latter respondents, we introduce another indicator for no industry or occupation code (see Garcia and Cowan, 2022).

Table C7 compares our baseline results (on the left) with estimates of equations (1)-(2) that include industry and occupation controls (on the right). With the added regressors, mothers' labor supply response is now estimated to be indistinguishable from zero. Estimates of fathers' labor supply response are less sensitive to the new controls, but even here the size of the coefficient is nearly halved in the specification with month-by-parental status effects.

One reason that the impact of industry and occupation controls varies by gender is that participation is a more active margin among women. Hence, the absence of industry and occupation codes is a more important predictor of women's hours and leaves little else for in-person shares to account for. To illustrate, we find that nearly 80 percent of nonemployed (zero-hours) women in the Outgoing Rotation Groups do *not* report an industry and occupation, whereas the analogue among men is 65 percent.

C.5 Measures of in-person share

This section reports on alternatives to the SafeGraph-based measure of the in-person share used in most of the main text. The first part of this section provides details on data from Burbio and the COVID-19 School Data Hub (CSDH) and presents additional results based on these sources. The second part examines alternative estimates from Parolin and Lee's analysis of SafeGraph data.

Table C7: Industry and Occupation Controls

	Baseline controls			With industry and occupation		
Coefficient	Women					
In-person × kids, ψ	0.582*	2.360***	-0.040	-0.165	0.526	0.303
	[0.304]	[0.634]	[0.672]	[0.202]	[0.358]	[0.438]
Number of obs.	447,899	447,899	447,899	447,881	447,881	447,881
Coefficient	Men					
In-person × kids, ψ	0.567*	1.888***	-0.051	0.577***	1.082***	-0.051
	[0.314]	[0.645]	[0.706]	[0.215]	[0.378]	[0.553]
Number of obs.	432,856	432,856	432,856	432,775	432,775	432,775
Month × parent	No	Yes	Yes	No	Yes	Yes
Area × parent	No	No	Yes	No	No	Yes

Note: The left panel reports estimates based on equations (1)-(2). The right panel adds controls for industry and occupation. Standard errors are clustered at the geographic area level. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

Alternatives to SafeGraph. Both Burbio and the CSDH are organized around three instruction formats: "in-person", "virtual", or "hybrid". For 46 states, CSDH classifies each school (or district if school-level data was not reported) under one of these formats. Enrollment figures are also provided, which enables us to calculate county-wide participation in each mode. In total, these states span almost 3,000 counties. Burbio directly reports county-level estimates of enrollment by format for just under 600 counties.

The in-person share encompasses both the self-reported "in-person" option as well as the in-person content of "hybrid" instruction. However, neither dataset specifies how many days per week students are in person under the "hybrid" format. To fill this gap, we rely on state-level tabulations by the Institute for Education Sciences (IES) at the U.S. Department of Education and derived from the 2021 National Assessment of Educational Progress Monthly School Survey. The survey was administered in each of the first five months of 2021, and results were published for 37 states (for which sufficient data was reported) as well as for the four Census regions. Since the survey was not fielded in 2020, we assume responses for January 2021 applied to earlier months of the school year.[8]

Based on IES figures, an in-person share of hybrid instruction is imputed as follows. Among schools in each state that report a "hybrid" format, the IES calculates the share for which the number of in-person days per week was (a) one to two, (b) three, or (c) four to five. We use these reports to calculate the share of weekly instruction held on-site under a "hybrid" format, where a two-day per week schedule is chosen to represent bin (a) and a four-day schedule represents bin (c). For states that did not participate in the survey, we substitute an estimate based on analogous tabulations for the Census region of the state. The final estimate is assigned to each of a state's counties. A county's overall in-person share is then computed as the sum of the share of the county's enrollment in the "in-person" format; and the share of the county's enrollment in the "hybrid" format scaled by our estimate of the share of hybrid instruction held on-site.

While we see a value in leveraging all the information available from IES, we have confirmed that simpler treatments of "hybrid" instruction yield similar results. For example, abstract from any variation in

[8] These data may be downloaded from `https://ies.ed.gov/schoolsurvey/mss-dashboard/`.

the in-person content of hybrid instruction and instead assign to all counties in all months the same on-site share of weekly instruction under a hybrid format. This share is the national mean in the 2020-21 school year and equal to 0.6, e.g., three days per week. We have found that estimates based on the latter are very similar to results in the main text (Table 4).[9] This conclusion suggests that variation in the overall in-person share is dominated by differences in the take-up of the three basic modes (in-person, hybrid, or fully virtual).

Next, we extend the analysis of CSDH data by investigating a subsample of states for which there is more refined enrollment data. Recall that our analysis in the main text draws on categorical data, e.g., the enrollment of a school or district *as a whole* is taken to be in one of the three formats. For 26 states, CSDH also report the distribution of enrollment *within* school (or district). These enrollment data enable a more precise estimate of the in-person share. Table C8 reports results for this subset of states and compares them to estimates based on the categorical data for this same set of states. As a further point of reference, we report estimates from SafeGraph for this subsample. Estimates for Burbio hardly differ relative to those shown in the main text and are omitted to conserve space. We also omit results that include spatial (area-by-parental status) controls since these estimates are uniformly insignificant.

The table reveals three noteworthy results. The first two columns indicate that estimates off categorical data for this subsample are similar to (but estimated less precisely than) those based on the full sample of states. The next two columns refer to regressions based on the detailed enrollment data. Estimated maternal labor supply responses increase by 1-1.5 hours per week, consistent with the claim that estimates off categorical data are attenuated. However, this result may also reflect an idiosyncratic feature of *this subset* of states that is not apparent in the noisier categorical data. The final two columns of the table provide some evidence in favor of this latter interpretation. Within this set of states, maternal labor supply responses in SafeGraph also increase by the same magnitude as in the CSDH data (see Tables 1-2 for a comparison to the full sample). The responses of fathers increase in the SafeGraph data, too, but not in the CSDH data.

Table C8: Comparison of In-person Shares in CSDH

	CSDH Categories		CSDH Enrollment		SafeGraph	
Coefficient	Women					
In-person × kids, ψ	0.974	1.213	2.367**	2.606**	3.452***	3.984***
	[0.807]	[0.899]	[0.951]	[1.018]	[1.023]	[1.119]
Number of obs.	94,390	94,390	94,390	94,390	94,390	94,390
Coefficient	Men					
In-person × kids, ψ	0.046	0.213	0.0524	0.192	2.460**	2.894 **
	[0.837]	0.872	[0.932]	[0.946]	[1.039]	[1.088]
Number of obs.	90,697	90,697	90,697	90,697	90,697	90,697
Month × parent	No	Yes	No	Yes	No	Yes

Note: In the first two columns, in-person shares are derived from categorical data in CSDH, e.g., a school or district is "in-person" or "hybrid". In the third and fourth columns, in-person shares are calculated from CSDH data on the enrollment distribution by instruction format within school or district. The final two columns are based on SafeGraph data. Throughout, the sample consists of local areas for which the detailed enrollment data is available. The time period is September 2020-May 2021. Standard errors are clustered at the geographic area level. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

[9] For instance, conditional on time-by-parental status controls, Table 4 reports $\psi = 1.423$ for women. The analogous result based on this simpler treatment of hybrid instruction is $\psi = 1.379$, which is also statistically significant.

Alternative SafeGraph measures. Table 5 in the main text illustrated how parental labor supply responses vary by age of the eldest child. For this exercise, we used our baseline in-person share from Parolin and Lee (2021), which measures on-site activity across *all* grades. This choice ensures a consistency with other regressions but fails to tailor the measurement of in-person shares to the age range of the children.

Therefore, Table C9 re-runs the regressions with in-person shares that apply specifically to the student ages in question. For ages 5-9, we use Parolin and Lee's estimate of the in-person share for elementary schools. Parolin and Lee also report a (single) estimate for all other grades. Absent a better alternative, we take the latter as the in-person share for middle schools and, separately, for high schools. Therefore, we adopt this measure for ages 14-17. Finally, for ages 5-13, we take an enrollment-weighted average of in-person shares in the elementary and non-elementary groups but where the weight on the latter is *middle* school enrollment. Data on enrollment are from the National Center for Education Statistics.[10]

These alternative in-person shares do not alter any of the conclusions drawn from Table 5. To conserve space, we present results in Table C9 only for the specification with month-by-parental status controls. Relative to Table 5, the point estimates in Table C9 are slightly larger for samples with younger kids (ages 5-13 and 5-9) and somewhat smaller for samples with older children (ages 14-17). Overall, though, these results buttress the conclusions in the main text.

Table C9: Student-age-specific In-person Share Estimates

	Ages 14-17	Ages 5-13	Ages 5-9
Coefficient		Women	
In-person × kids, ψ	1.129	2.592***	2.831***
	[0.853]	[0.709]	[0.888]
Number of obs.	325,420	405,165	358,146
		Men	
In-person × kids, ψ	0.953	2.201***	3.033***
	[0.911]	[0.696]	[0.850]
Number of obs.	332,848	399,384	362,062

Note: In columns marked "Ages 14-17", mothers are included only if their eldest child is between ages 14-17. The columns, "Ages 5-13" and "Ages 5-9", are defined analogously. The period is calendar years 2020 and 2021 but with June-August omitted. Parental status-by-month effects are included throughout. See text for definitions of the in-person shares. Standard errors are clustered at the geographic area level. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

C.6 Comparison to Garcia and Cowan (2022)

Table C10 compares our estimation to that of Garcia and Cowan (2022). To start, we apply the latter's specification. Our estimates are reported in the first column on the left. Results for women and unmarried men are near replicas of theirs (see their Tables 5-6). For married men, our estimate is somewhat higher: we recover a coefficient of 1.961 whereas Garcia and Cowan find 1.749.

[10] See the Elementary/Secondary Information System (ElSi) at `https://nces.ed.gov/ccd/elsi/`.

Table C10: Reconciliation with Garcia and Cowan (2022)

Coefficient	Married women						
In-person × kids, ψ	0.803**	0.877***	0.967***	1.245***	2.037***	0.812**	2.256***
	[0.332]	[0.284]	[0.291]	[0.195]	[0.310]	[0.362]	[0.788]
Number of obs.	168,860	169,018	152,595	389,194	389,201	242,743	242,743
Coefficient	Unmarried women						
In-person × kids, ψ	1.672***	2.093***	2.302***	1.865***	2.359***	0.957*	2.591**
	[0.447]	[0.449]	[0.452]	[0.318]	[0.517]	[0.543]	[1.032]
Number of obs.	164,043	164,145	148,221	339,543	339,557	205,156	205,156
Coefficient	Married men						
In-person × kids, ψ	1.961***	2.205***	2.326***	2.197***	2.377***	1.276***	1.824***
	[0.325]	[0.294]	[0.299]	[0.219]	[0.305]	[0.338]	[0.661]
Number of obs.	167,545	169,280	152,837	388,503	388,547	223,471	223,471
Coefficient	Unmarried men						
In-person × kids, ψ	1.926**	2.028**	2.257***	2.470***	2.415***	1.231*	1.657
	[0.792]	[0.794]	[0.776]	[0.535]	[0.664]	[0.675]	[1.455]
Number of obs.	132,409	132,784	120,262	282,811	282,856	209,385	209,385
Weighted	Yes	No	No	No	No	No	No
Includes Aug.	Yes	Yes	No	No	No	No	No
Only CPS counties	Yes	Yes	Yes	No	No	No	No
Industry and occupation	Yes	Yes	Yes	Yes	No	No	No
Age range (of adults)	21+	21+	21+	21+	21+	21-59	21-59
Time × parent effects	No	No	No	No	No	No	Yes

Note: The outcome variable is weekly hours worked. The period is calendar years 2020 and 2021 with only June-July omitted if "Includes Aug." is "Yes"; otherwise, August is also omitted. The "CPS counties" are FIPS counties identified in the CPS; if this row is "No", then all local areas are in the sample. See the text for a discussion of the other sample selection rules and specifications. Standard errors are clustered at the geographic area level. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

The regressions in the first column are weighted by CPS sample weights, as in Garcia and Cowan. We opted not to weight in the main text, though. Therefore, our next step is to drop the weights. A comparison across the first (weighted) and second (unweighted) columns indicates that weighting implies somewhat smaller estimates, but the impact is rather modest. For married women and unmarried men in particular, weighting makes little difference. For married men and unmarried women, weighting has a more noticeable impact—the hours response increases by around 0.25 to 0.40 weekly hours, respectively. Still, one could not reject that the weighted and unweighted estimates are equivalent.

As we proceed further to the right of the table, we make several more adjustments to Garcia and Cowan's sample and specification. These modifications are to (i) drop the month of August from the sample, as many school districts (in the Northeast in particular) do not open until September (DeSilver, 2023); (ii) add local areas in addition to the counties identified in the CPS; (iii) remove the industry and occupation controls; and (iv) restrict the age range (of adult respondents) to 21-59. The first three adjustments tend to *elevate* the hours responses relative to results in Garcia and Cowan, most notably for married women. As noted earlier, though, the final adjustment—the restriction to ages 21-59—implies estimates that in most cases are at least 50 percent *smaller*. The final column on the far right introduces the month-by-parental status effects. This specification recapitulates what appears in Tables 6 and 7 in the main text. With these controls, the estimates are higher and, in the case of men, nearly on par with Garcia and Cowan's original results. Our final estimates for women are higher, however, than in the latter's paper.

C.7 Married couples' hours worked

We close by examining the joint hours response of married couples, that is, the response of total hours worked of the couple to a shift in instruction format. However, the CPS does not explicitly identify each respondent's married partner. Therefore, to ensure that two self-identified married respondents within a family unit are indeed a couple, each must report that his/her married partner is "present". This restriction removes 3 percent of married respondents in our original sample. We retain any other married couple in which at least one member is part of our original sample.

Table C11 presents the results. The top panel reports the main estimates of interest, namely, the response of the couple's total hours to a shift from fully virtual to fully in-person. As a point of reference, the bottom two panels present separate estimates for the mothers and fathers who are in married couples and in our original sample.[11] The regression specification for the bottom two panels is identical to that used in Section 3, e.g., equation (1) augmented with month-by-parental-status effects. In the top panel, we extend this specification to include a full set of demographic controls for *each* member of the couple.

Table C11: Married Couple Hours of Work

	All married	Married College	Married Noncollege
Coefficient		Married couple	
In-person × kids, ψ	3.651***	2.793*	5.448***
	[1.045]	[1.654]	[1.642]
Number of obs.	245,179	74,526	112,326
Coefficient		Women in married couple	
In-person × kids, ψ	2.157***	2.768**	2.729**
	[0.803]	[1.380]	[1.067]
Number of obs.	235,808	72,649	106,949
Coefficient		Men in married couple	
In-person × kids, ψ	1.714**	0.854	2.627**
	[0.668]	[0.958]	[1.079]
Number of obs.	216,753	68,130	97,639

Note: A married college (married noncollege) couple is one in which both members are college (noncollege) graduates. The unit of analysis in the top panel is the couple. In the other panels, it is the individual married respondent. The period is calendar years 2020 and 2021 but with June-August omitted. Parental status-by-month fixed effects are included throughout. Standard errors are clustered at the geographic area level. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

[11] Of those in married couples, over 92 percent are in our original (e.g., ages 21-59) sample. The broader sample of individuals in couples accounts, in large part, for why the number of couples in the top panel of Table C11 exceeds the number of individual mothers or fathers in the bottom panels. In addition, the presence of same-sex couples implies a discrepancy between the number of couples, on the one hand, and the number of mothers and fathers on the other.

Our takeaways from Table C11 are as follows. First, a shift from fully virtual to in-person implies an increase in couples' total hours of work of 3.7 per week. This result is remarkably similar to, if slightly smaller than, the sum of responses reported in the main text for married men and, separately, married women (4.1 hours per week—see Table 7). Second, the *distribution* of the hours responses within the household varies notably by education. The joint hours response of married, college educated couples largely reflects labor supply adjustments of mothers. By contrast, in households with two noncollege educated spouses, the labor supply adjustments are virtually identical.

A simple framework for interpreting results in Table C11 connects differences in labor supply behavior to differences in earnings opportunities. For instance, responding to a local school closure, a father may continue working if his earnings are highest, with his spouse allocating more time to childcare. Conversely, if parents' earnings are similar, they may take up childcare to a similar extent. Thus, the pattern in Table C11 can potentially arise if a father's relative earnings (within the household) are *increasing* in his schooling, e.g., if a college-educated father is more likely to have earnings exceeding those of his spouse.

To pursue this point further, we draw on weekly earnings records from the CPS Outgoing Rotation Groups. The data are from 2019 and, therefore, capture the situation facing parents prior to when on-site instruction was suspended. Earnings of the non-employed are set to zero. Interestingly, we find that fathers' relative earnings are nearly *independent* of schooling. Two moments of the data demonstrate this point. First, roughly 70 percent of fathers earn more than their spouses regardless of college attainment. These figures echo results in Winkler et al. (2005), who used annual earnings from the CPS March Supplement. Second, we compute the father's earnings premium as the difference between his and his spouse's earnings relative to the couple's average earnings. The average premium varies little with schooling, falling between 70 (noncollege) and 73.5 (college) percent. These moments reflect the tendency of fathers to have spouses with the same level of schooling, e.g., one partner's college premia is balanced by the other.[12]

These findings challenge an explanation of the gender gradient in Table C11 based on differences in returns to work. One caveat, though, is that current earnings do not fully reflect returns on market time. Since returns to experience may be somewhat higher for men (see Munasinghe, Reif, and Henriques, 2008), a household may select the father for full-time work even if parents' current earnings are similar. Alternatively, our estimates may point to differences in preferences for and/or norms around childcare. It is an open question, though, why these differences would pertain only to college graduates.

D. Additional results from the ATUS

This appendix reports additional estimates from the American Time Use Survey (ATUS). To start, we examine each major activity category in Table D1. On the whole, we do not detect a significant response of time use in any activity to instruction format. The one exception is the reduction in fathers' time spent in childcare when in-person instruction returns. Recall, however, that the table pertains only to time use in the primary activity. These results do not capture time spent on the joint performance of work (the primary activity) and childcare (secondary). We turn to this issue next.

In Table D2, we report on time spent working from home *while* supervising children. (The response of total telework hours remains insignificant.) This margin was more active among women, especially among college-educated women. For instance, the return of on-site instruction reduced time in this activity

[12] Nearly 80 percent of married fathers with a college degree have college-educated spouses. Similarly, almost three-quarters of noncollege-educated fathers have spouses with no more than a high school degree.

by over 8 hours per week among women with college degrees. Responses of college educated men and noncollege educated women were insignificant.

Table D3 addresses nonparental care supplied by over-age-60 respondents. The coefficient measures the impact of a shift to in-person instruction on their weekly hours spent supervising a child under 18. The table confirms that, in the unweighted regressions, the response of nonparental time was similar for college and noncollege groups. As we will see, there is an educational gradient in the weighted OLS results.

Table D1: Instruction Format and Time Use Across Major Activity Categories

Coefficient	Work	Leisure	Childcare	Home prod.	Commute
			All		
In-person share, β	0.565	-3.431	1.988**	0.504	0.564
	[3.973]	[2.717]	[0.888]	[2.054]	[0.562]
In-person × kids, ψ	1.839	0.166	-2.408	-2.621	0.239
	[4.256]	[2.816]	[1.946]	[1.824]	[0.533]
Number of obs.	3,278	3,278	3,278	3,278	3,278
			Men		
In-person share, β	6.284	-1.510	0.199	-1.147	0.320
	[5.906]	[4.280]	[0.893]	[2.463]	[0.934]
In-person × kids, ψ	-0.003	-6.030	-3.714**	0.114	1.016
	[6.699]	[4.731]	[1.753]	[2.876]	[1.048]
Number of obs.	1,476	1,476	1,476	1,476	1,476
			Women		
In-person share, β	-2.459	-2.566	3.438**	0.589	0.890
	[4.754]	[3.746]	[1.533]	[2.692]	[0.723]
In-person × kids, ψ	2.267	2.323	-0.936	-3.098	-0.221
	[5.088]	[3.243]	[2.277]	[2.618]	[0.577]
Number of obs.	1,701	1,701	1,701	1,701	1,701
			Noncollege		
In-person share, β	-1.657	-2.212	1.341	1.168	0.194
	[6.338]	[4.232]	[1.328]	[2.989]	[0.583]
In-person × kids, ψ	1.125	1.863	-3.590	0.795	-0.167
	[7.168]	[5.296]	[3.202]	[3.423]	[0.887]
Number of obs.	1,623	1,623	1,623	1,623	1,623
			College		
In-person share, β	6.540	-6.662*	2.122	-0.001	0.479
	[5.370]	[3.805]	[1.478]	[2.858]	[1.041]
In-person × kids, ψ	0.315	1.787	-2.054	-4.516	0.142
	[6.146]	[3.181]	[2.006]	[2.770]	[0.892]
Number of obs.	1,561	1,561	1,561	1,561	1,561

Note: Each column by panel is a separate regression, with the implied number of hours per week spent in each activity as the dependent variable. (Daily hours are multiplied by five.) The sample consists of observations between Monday and Friday. Relative to equation (1), the specification also includes fixed effects for days of the week and parent status×month. Standard errors are clustered at the geographic area level. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

Table D2: Working From Home While Caring For Children

	All	Men	Women
Coefficient		All	
In-person share, β	1.553*	2.115**	0.399
	[0.935]	[0.972]	[1.512]
In-person × kids, ψ	-5.897***	-4.080**	-6.320***
	[1.535]	[2.013]	[2.347]
Number of obs.	3,278	1,476	1,701
		Noncollege	
In-person share, β	-0.455	1.027	-1.226
	[0.703]	[0.988]	[1.370]
In-person × kids, ψ	-1.903	-4.070	0.649
	[2.290]	[3.722]	[3.484]
Number of obs.	1,623	722	806
		College	
In-person share, β	2.800*	2.301	1.909
	[1.566]	[2.046]	[2.477]
In-person × kids, ψ	-6.652**	-3.186	-8.241**
	[2.704]	[3.833]	[3.854]
Number of obs.	1,561	664	806

Note: Each column by panel is a separate regression, with the sample defined by the column header. The dependent variable is the implied number of hours per week where "work at home" is the primary activity and "childcare" is secondary. (Daily hours are multiplied by five.) Relative to equation (1), the specification includes fixed effects for days of the week and parent status×month. Standard errors are clustered at the geographic area level. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

Table D3: Nonparental Care

	All	Men	Women
Coefficient		All	
In-person share, β	-3.385*	-2.547	-3.878**
	[1.480]	[1.734]	[2.165]
Number of obs.	2,425	976	1,354
		Noncollege	
In-person share, β	-2.752	-1.301	-3.281
	[1.834]	[1.996]	[2.387]
Number of obs.	1,453	540	799
		College	
In-person share, β	-2.836	-1.966	-2.460
	[2.770]	[5.705]	[4.072]
Number of obs.	877	344	453

Note: Each column by panel is a separate regression, with the sample defined by the column header. The implied number of hours per week spent with other's children is the dependent variable. (Daily hours are multiplied by five.) Time spent with other's children includes any time spent with a person under 18 years old outside of market work. Relative to equation (1), the specification includes fixed effects for days of the week and parent status×month. Standard errors are clustered at the geographic area level. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

Lastly, we report regression results based on ATUS sample weights. Table D4 concerns telework. The table shows that, after a shift to in-person instruction, the weekly hours spent working from home while caring for children falls by 6.3 among college educated parents. The counterpart in the main text is 6.6 (Table 8), which shows that weighting makes little difference to the result. Table D5 reports on nonparental care. In the main text, we estimated a significant reduction of 3.4 weekly hours driven by changes in women's time use (Table 9). In the weighted results in Table D5, the point estimate is smaller and insignificant in the full sample. Rather, there is a significant response only among the noncollege educated.

Table D4: Weighted Estimates—Working From Home, Childcare, and Instruction Format

	Work	Work at home	Work at home + childcare
Coefficient		All	
In-person share, β	0.439	-4.318	1.093
	[4.903]	[3.009]	[0.657]
In-person × kids, ψ	-4.575	-2.585	-4.792***
	[5.184]	[4.693]	[1.609]
Number of obs.	6,622	6,622	6,622
		Noncollege	
In-person share, β	0.516	0.096	-0.164
	[7.004]	[3.381]	[0.386]
In-person × kids, ψ	-6.647	1.624	-1.069
	[7.896]	[4.705]	[1.963]
Number of obs.	3,371	3,371	3,371
		College	
In-person share, β	3.113	-5.406	1.176
	[5.993]	[5.004]	[1.040]
In-person × kids, ψ	0.300	-3.662	-6.308**
	[6.051]	[7.975]	[2.781]
Number of obs.	3,178	3,178	3,178

Note: Each column by panel is a separate regression, with the implied number of hours per week spent in each type of activity as the dependent variable. (Daily hours are multiplied by seven.) All days of the week are included, and ATUS sample weights are used. See Notes to Table D2 for more. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

Table D5: Weighted Estimates—Nonparental Care

	All	Men	Women
Coefficient		All	
In-person share, β	-1.753	-0.477	-2.961
	[2.082]	[2.695]	[2.535]
Number of obs.	4,848	1,983	2,787
		Noncollege	
In-person share, β	-4.261*	-0.401	-6.445**
	[2.287]	[2.626]	[3.105]
Number of obs.	2,945	1,106	1,725
		College	
In-person share, β	1.029	-0.681	1.957
	[3.734]	[6.511]	[4.491]
Number of obs.	1,817	765	952

Note: Each column is a separate regression, with the implied number of hours per week spent with other's children as the dependent variable. (Daily hours are multiplied by seven.) All days of the week are included, and ATUS sample weights are used. See Notes to Table D3 for more. *** indicates a p-value less than 0.01; ** a p-value between 0.01 and 0.05; and * a p-value between 0.05 and 0.10.

E. Derivations and proofs

This appendix derives results for the model with nonparental care introduced in Section 5. To start, we restate the parent's problem:

$$\max_{c,l,m,x} \alpha \ln c + \beta \ln l + (1-\alpha-\beta) \ln q$$

subject to the child development technology,

$$q = g^{\gamma} \left[\left(\mu^{1-\psi} m^{\psi} + (1-\mu)^{1-\psi} x^{\psi} \right)^{1/\psi} \right]^{1-\gamma} \tag{E.1}$$

with $\psi \leq 1$. The time constraints of the adult and child are, respectively:

$$1 = l + m + n, \quad \text{and}$$

$$1 = g + m + x.$$

In addition, the parent faces a budget constraint of the form,

$$c + px + \text{savings} = wn + \text{asset income}\,.$$

For our purposes, we do not need to specify the full asset allocation problem.

The (intra-temporal) first order conditions may be condensed to two expressions, one for market time, n, and the other for nonparental time, x. The optimal choice of market time, originally stated in the main text, is repeated here,

$$n = g + x - \frac{\beta}{\lambda w}. \tag{E.2}$$

The demand for nonparental time satisfies

$$(1-\alpha-\beta)(1-\gamma)\frac{\mu^{1-\psi}(1-g-x)^{\psi-1}-(1-\mu)^{1-\psi}x^{\psi-1}}{\mu^{1-\psi}(1-g-x)^{\psi}+(1-\mu)^{1-\psi}x^{\psi}}=\lambda(w-p), \tag{E.3}$$

where λ is the marginal value of wealth. Note that if $w>p$, then the optimal choice of x requires

$$1<\frac{\mu}{1-\mu}\cdot\frac{x}{1-g-x}.$$

In what follows, we let

$$\nu\equiv\frac{\mu}{1-\mu}\ \text{ and }\ \xi\equiv\frac{x}{m}=\frac{x}{1-g-x}.$$

Now take logs of equation (E.3) and totally differentiate with respect to x and g. The comparative static may be expressed as

$$\frac{\mathrm{d}x}{\mathrm{d}g}=-\frac{1}{1+z(\xi;\psi)}, \tag{E.4}$$

where we have defined

$$z(\xi;\psi)\equiv\frac{(\nu\xi)^{\psi-1}+(1-\psi)\xi^{-1}-\psi}{(\nu\xi)^{1-\psi}+(1-\psi)\xi-\psi}. \tag{E.5}$$

The sign of this comparative static is unambiguous, as shown next.

Lemma 1. *Nonparental care, x, declines in publicly provided supervision,* $\mathrm{d}x/\mathrm{d}g<0$.

Proof. We must establish that $1+z>0$. Recall that $\nu\xi>1$ for an interior solution with $w>p$. It follows that the denominator in z must be positive: $(\nu\xi)^{1-\psi}+(1-\psi)\xi-\psi>(1-\psi)(1+\xi)>0$. Therefore, $1+z>0$ if

$$(\nu\xi)^{\psi-1}+(\nu\xi)^{1-\psi}+(1-\psi)(\xi+\xi^{-1})-2\psi>0. \tag{E.6}$$

Since $\xi\in\mathbb{R}^{+}$, the term $\xi+\xi^{-1}$ attains a minimum of 2 at $\xi=1$. Likewise, define $\rho\equiv(\nu\xi)^{1-\psi}>1$ and note that, by the same logic, $\rho+\rho^{-1}$ is no smaller than 2. Hence, the left side of equation (E.6) has a minimum of $2+2(1-\psi)-2\psi=4(1-\psi)>0$. This confirms that $\mathrm{d}x/\mathrm{d}g<0$. ∎

The comparative static for market time follows from equations (E.2) and (E.4),

$$\frac{\mathrm{d}n}{\mathrm{d}g}=\frac{z(\xi;\psi)}{1+z(\xi;\psi)}. \tag{E.7}$$

The sign of this comparative static depends on ψ. From equation (E.5), $z>0$ for any $\psi\leq 0$. However, for ψ sufficiently near one, $z=(\rho^{-1}-\psi)/(\rho-\psi)<0$. The behavior of $\mathrm{d}n/\mathrm{d}g$ inherits these properties.

More generally, we establish that z, and by extension $\mathrm{d}n/\mathrm{d}g$, declines monotonically in ψ given an initial solution ξ that satisfies equation $\nu\xi>1$. That is, we can characterize the map from ψ to z local to an initial optimum. This approach to comparative statics on ψ is akin to the "normalization" advocated by La Grandville (1989) and Klump and La Grandville (2000) when one works with CES functions (see Cantore and Levine, 2014, on this point). To perturb ψ but hold ξ fixed, the share parameter, $\nu\equiv\mu/(1-\mu)$, is adjusted as needed.

Proposition 1. *Given a solution ξ that satisfies $\nu\xi > 1$, the comparative static of market time* $\mathrm{d}n/\mathrm{d}g$ *declines monotonically in ψ and crosses zero once at a threshold $\hat{\psi} \in (0,1)$.*

Proof. As a first step, we determine how ν must be adjusted so that any initial optimum ξ still holds after ψ is perturbed. Recall $\rho \equiv (\nu\xi)^{1-\psi} > 1$ and rewrite equation (E.3) as

$$\frac{\rho - 1}{\rho + (x/(1-g-x))} = (1-g-x) \cdot \frac{\lambda(w-p)}{(1-\alpha-\beta)(1-\gamma)}.$$

This expression indicates that to hold x—and by extension $\xi \equiv x/(1-g-x)$—fixed at an initial optimum, it is necessary to adjust ν such that ρ does not change, e.g., $d\rho = 0$. Therefore, a perturbation $\mathrm{d}\psi$ requires an adjustment $\mathrm{d}\nu = \mathrm{d}\psi \cdot \nu \ln(\xi\nu)/(1-\psi)$. Now totally differentiate equation (E.5) with respect to ν and ψ subject to the latter restriction and given an initial optimum ξ_0 and a $\rho = \rho_0 > 1$. The comparative static is

$$\left.\frac{\mathrm{d}z}{\mathrm{d}\psi}\right|_{\xi=\xi_0} = -\frac{(1+\xi_0^{-1})(\rho_0 - 1) + (1+\xi_0)(1-\rho_0^{-1})}{(\rho_0 + (1-\psi)\xi_0 - \psi)^2} < 0.$$

Since $\mathrm{d}n/\mathrm{d}g$ is a monotone function of z, the former also declines in ψ. Moreover, since $z > 0$ for any $\psi \leq 0$ but turns negative as $\psi \to 1$, the point $\hat{\psi}$ at which $\mathrm{d}n/\mathrm{d}g$ crosses zero must be strictly between zero and one. ■

The connection between $\mathrm{d}n/\mathrm{d}g$ and ψ enables us to draw inferences about the latter given estimates of the former. Specifically, we describe in the main text how to bound the range of ψs consistent with a sufficiently small $\mathrm{d}n/\mathrm{d}g$. This result is formally stated below.

Corollary 1. *For dn/dg sufficiently small, ψ is bounded from below such that $\psi > (1+\xi)^{-1}$.*

Proof. Consider first the special case where $\mathrm{d}n/\mathrm{d}g = 0$. Fix ξ and ρ. Equations (E.5) and (E.7) imply that ψ satisfies

$$\psi = \frac{\rho^{-1} + \xi^{-1}}{1 + \xi^{-1}} > \frac{1}{1+\xi}, \tag{E.8}$$

which confirms that the bound obtains at $\mathrm{d}n/\mathrm{d}g = 0$. More generally, let the comparative static take the value $\delta/(1+\delta) > 0$ with $\delta \in \mathbb{R}^+$. It follows that $z = \delta$, which mean ψ satisfies

$$\psi = \frac{\rho^{-1} + \xi^{-1} - \delta(\rho + \xi)}{1 + \xi^{-1} - \delta(1+\xi)}.$$

To a first order around $\delta = 0$, this is given by

$$\psi = \frac{1}{1+\xi} + \frac{\rho^{-1}}{1+\xi^{-1}} - \frac{(1+\xi)(1-\rho^{-1}) + (1+\xi^{-1})(\rho - 1)}{(1+\xi^{-1})^2}\delta.$$

Consider the difference between the two terms on the right. The bound in equation (E.8) will still apply if this difference is positive, which will obtain for all δ such that

$$\frac{1}{(\rho+\xi)(\rho-1)} > \delta. \tag{E.9}$$

In this sense, the bound applies for δ sufficiently small. ■

To illustrate the result in equation (E.9), fix $\xi = 1.3$ as in the main text and consider $\psi = 2/3$. Further, suppose $\mu = 4/5$ is consistent with the choice of $\xi = 1.3$ (given $\psi = 2/3$, wage rate w, care price p, and so on). Therefore, $\nu \equiv \mu/(1-\mu) = 4$ and $\rho = 1.7325$. Equation (E.9) then requires $\delta < 0.44$, or $\mathrm{d}n/\mathrm{d}g = \delta/(1+\delta) < 0.3$, which is easily satisfied. The upper bound on δ *rises* at lower values of μ.

Next, the main text observes that the response of market time to variation in the price of nonparental care, p, also hinges on the value of ψ. To illustrate this point, consider a *temporary* decrease in p as would be implied by the subsidies provided in the pandemic period. Of course, a *permanent* decrease will have an income effect, which will mitigate the impact of the price on labor supply shown below. Total differentiation of equation (E.3) yields

$$\left\{(1-\psi)\frac{\xi\rho+1}{\rho-1} + \psi\xi\frac{\rho-1}{\rho+\xi}\right\}\mathrm{d}\ln x = -\frac{p}{w-p}\mathrm{d}\ln p, \tag{E.10}$$

where ρ is defined as before. It is immediate that the left side of this expression is positive for any $\psi \in [0,1]$. It remains positive for $\psi < 0$ if $(1-\psi)$ multiplies a larger number than ψ. This is the case since

$$\frac{\xi\rho+1}{\rho-1} > \frac{\xi\rho-\xi}{\rho+\xi}.$$

Now collect terms on the left side of equation (E.10) that involve ψ to rewrite this expression as

$$\frac{\mathrm{d}\ln x}{\mathrm{d}\ln p} = -\frac{p}{w-p} \times \frac{(\rho-1)(\rho+\xi)}{(\xi\rho+1)(\rho+\xi)-\psi\rho(1+\xi)^2} < 0.$$

The magnitude of the term on the right increases in ψ for given ξ and ρ, which confirms that a lower price stimulates more nonparental time if the two forms of care are more substitutable. The same logic applies to the comparative static for market time, which is given by

$$\frac{\mathrm{d}\ln n}{\mathrm{d}\ln p} = \frac{x}{n} \times \frac{\mathrm{d}\ln x}{\mathrm{d}\ln p} = -\frac{x}{n} \times \frac{p}{w-p} \times \frac{(\rho-1)(\rho+\xi)}{(\xi\rho+1)(\rho+\xi)-\psi\rho(1+\xi)^2}. \tag{E.11}$$

Finally, the market time response to a temporary wage increase follows the same steps as those above. Specifically, one merely needs to swap the term $-\frac{p}{w-p} < 0$ for $\frac{w}{w-p} > 0$ in equation (E.11) to express the comparative static.

Additional References

Almagro, Milena and Angelo Orane-Hutchinson. 2020. "JUE Insight: The Determinants of the Differential Exposure to COVID-19 in New York City and Their Evolution Over Time," *Journal of Urban Economics*, 127.

Cantore, Cristiano and Paul Levine. 2012. "Getting Normalization Right: Dealing with 'Dimensional Constants' in Macroeconomics." *Journal of Economic Dynamics and Control*, 36(12): 1931–1949.

Carozzi, Felipe, Sandro Provenzano, and Sefi Roth. 2022. "Urban Density and COVID-19: Understanding the US Experience." *Annals of Regional Science.*

Case, Anne and Angus Deaton. 2021. "Mortality Rates by College Degree Before and During COVID-19," NBER Working Paper 29328.

De La Grandville, Olivier. 1989. "In Quest of the Slutsky Diamond." *American Economic Review*, 79(3): 468–481.

DeSilver, Drew. 2023. "`Back to school' Means Anytime from Late July to After Labor Day, Depending on Where in the U.S. You Live," Pew Research Center. Retrieved August 2023 from `https://pewrsr.ch/3QVWPKT`.

Kimmel, Jean. 1998. "Child Care Costs as a Barrier to Employment for Single and Married Mothers," *Review of Economics and Statistics*, 80(2): 287–299.

Klump, Rainer, and Olivier de La Grandville. 2000. "Economic Growth and the Elasticity of Substitution: Two Theorems and Some Suggestions." *American Economic Review,* 90(1): 282–291.

Manson Steven, Jonathan Schroeder, David Van Riper, Tracy Kugler, and Steven Ruggles. 2022. IPUMS National Historical Geographic Information System: Version 17.0 [dataset]. Minneapolis, MN: IPUMS.

Munasinghe, Lalith, Tania Reif, and Alice Henriques. 2008. "Gender Gap In Wage Returns to Job Tenure and Experience." *Labour Economics*, 15(6): 1296–1316.

Winkler, Anne E., Timothy D. McBride, and Courtney Andrews. 2005. "Wives Who Outearn Their Husbands: A Transitory or Persistent Phenomenon for Couples?" *Demography*, 42(3): 523–535.