FEDERAL RESERVE BANK
PHILADELPHIA

Working Papers
RESEARCH DEPARTMENT

School Closures, Parental Labor Supply, and Time Use

Enghin Atalay
Federal Reserve Bank of Philadelphia Research Department

Ryan Kobler
Federal Reserve Bank of Philadelphia Research Department

Ryan Michaels
Federal Reserve Bank of Philadelphia Research Department

WP 25-33

PUBLISHED
October 2025

ISSN: 1962-5361

DOI: https://doi.org/10.21799/frbp.wp.2025.33

SCHOOL CLOSURES, PARENTAL LABOR SUPPLY, AND TIME USE

Enghin Atalay Ryan Kobler Ryan Michaels

October 2025

Abstract

This paper re-examines the response of parental labor supply to the pandemic-era suspension of in-person instruction. The effect of school closures is undetectable after controlling comprehensively for unobserved heterogeneity. Even excluding such controls, a shift from fully virtual to in-person implies an increase in weekly hours worked of just 2 to 2.5. These estimates are used to inform a simple model of the household in which access to telework and nonparental care mitigate the labor supply impact of school closures. Time use data suggest telework and nonparental care indeed helped some parents balance work and childcare during the pandemic.

JEL codes: J21, J22, J48

Keywords: Market work, family economics, childcare, pandemic, in-person learning

---

Atalay, Kobler, and Michaels: Research Department, Federal Reserve Bank of Philadelphia. We thank Mike Dotsey, who stimulated our interest in the topic, and Mike Zabek for helpful discussions. We also thank participants at Midwest Macro (Richmond, 2024) and ESPE (Rotterdam, 2024). 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 suspension of in-person instruction was widely expected to upend working parents' careers (Goldin, 2022). However, initial analyses did not point to a dramatic change in parents' working time (Goldin, 2022; Furman et al., 2021).

Prompted by these findings, we first re-evaluate the impacts of school closures on parental labor supply. We consider a variety of specifications that, taken together, suggest the true response lies within a fairly tight range that includes zero. This raises interesting questions for the economics of time use, which we then examine. How did parents ease the trade-off between market work and childcare? On what margins, beyond labor supply, did they adjust? And what do these decisions imply about the preferences, technologies, and constraints shaping parents' time allocation decisions?

As a first step, we revisit evidence on the effect of remote instruction on parental labor supply. Following leading work by Garcia and Cowan (2024) and Hansen et al. (2024), we link adults' working time to the local schooling mode. As detailed in Section 1, the in-person share of instruction time is based on visits to school campuses as captured by SafeGraph's mobile phone location data (Parolin and Lee, 2021). When aggregated to a county or larger unit, these estimates can be matched to individuals (in that local area) in the Current Population Survey (CPS).

The merged SafeGraph-CPS dataset is the main source for our labor supply analysis. We estimate regressions that relate individual hours of work to the local in-person share of instruction. The potential endogeneity of school policy complicates the interpretation of this estimate, however. For example, parents' labor supply and school policy may be jointly shaped by local institutions and preferences.

While there is no "silver bullet" for this endogeneity problem, we present evidence based on a range of specifications that the causal effect is likely to lie within a reasonably narrow set of estimates. Our most parsimonious model follows the standard practice in this literature, which is to use childless adults as a control group for parents (Garcia and Cowan, 2024; Heggeness and Suri, 2021). This leverages within-area variation in working time across adults with and without children, effectively differencing out area-wide factors. We also consider richer specifications that include fixed effects which interact parental status with time and with local area. These additional fixed effects are motivated by the idea that, insofar as parents and nonparents have different preferences and market opportunities, this heterogeneity may vary across space as well as over time in the pandemic period.

Section 3 presents our main labor supply results. In the most parsimonious specification, a switch from virtual to in-person instruction lifts parents' hours by 0.5 per week relative to those of childless adults. With parental status-by-time effects included, the coefficient jumps to around 2.5 hours per week for mothers and 1.6 for fathers. We trace the source of this difference to what happened in the first half of 2020. We suggest that the original estimate of the effect on parental hours may be depressed because widespread school closures at that time coincided with exceptional labor market turbulence for childless adults.

The introduction of parental status-by-area effects has the opposite impact: it *eliminates* any association between in-person shares and parents' relative hours (that is, hours relative to those of adults with no kids). A null effect may arise if school policies are correlated with long-standing spatial differences in parental working time. However, the addition of more fixed effects also risks saturating the model. To arbitrate this issue, Section 3 conducts a placebo test: are school policies correlated with *pre*-COVID labor market outcomes? Indeed, higher average pandemic-era in-person shares predict higher parents' relative hours worked in the pre-pandemic period, particularly for mothers.

The placebo results suggest that specifications without parent-by-area effects are likely to yield upwardly biased estimates. This is notable since an hours response of around two is modest. The latter may, however, mask larger shifts by some parents. To probe how high such estimates may go, we extend our analysis without parent-by-area effects to various demographic groups.

We highlight three results on the heterogeneity of labor supply responses. First, estimates are similar across levels of educational attainment apart from college-educated fathers, who are essentially unresponsive to the in-person share. Second, among parents of younger school-age children (i.e., with children aged 5-9), hours adjust by as much as three per week. Third, labor supply responses vary little by marital status but do vary within the unmarried. Labor supply is relatively elastic among lone-adult parents—weekly hours rise by as much as four when in-person instruction is reinstated—but unresponsive among the unmarried in co-residential arrangements.

To take stock, we see a labor supply response of 2-4 weekly hours as the upper end of any plausible range of estimates. This figure is a small fraction of the roughly 30 hours of on-site time at reopened schools. This observation suggests that parents must have adjusted time use on other margins so as to both attend to children and supply labor.

To this end, we next report on several results from the American Time Use Survey (ATUS). First, there was in fact little adjustment in leisure, market work, or 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 in the pandemic period. Respondents over age 60—a group likely to include many grandparents—allocated up to 4 more hours per week to the care of *others'* children when in-person instruction was suspended. This response was observed only among those with no college degree.[1] Because of the small sample size of the ATUS, these estimates are subject to considerable uncertainty. Still, the results point to two promising explanations of the labor supply findings.

In Section 5, we view these results through the lens of simple models of parental time allocation. To start, we consider a baseline with no telework or nonparental care. Following Berlinski et al. (2024), a parent in the model values consumption, leisure time, and child development. In this setting, a child's development is a function of (only) 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. We show that a decline in publicly provided supervision leads the parent to substitute time toward childcare and yields a reduction in labor supply that is *at least* four times larger than our upwardly biased estimates (see above). In this sense, the regression estimates appear to be remarkably small.

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 parameter, and we derive the mapping from this parameter to the labor supply response. Second, noting that

[1] We view grandparents' educational attainment as the best available proxy for that of the parent. Unfortunately, the ATUS does not report the identity or the educational attainment of the parent of the child who received care from the over-age-60 respondent. We return to this issue later in Section 4.

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., 2024). This section concludes by highlighting the broader implications of this substitutability for public policy and cyclical hours dynamics.

Related research. Our paper intersects with several strands of research. First, our analysis of CPS data contributes to the literature on labor supply in the pandemic period. Our placebo test probes for the endogeneity of school policy and frames results from standard, two-way fixed effects models as upwardly biased estimates of the true effect. We see this approach as a complement to some earlier efforts. For instance, Hansen et al. (2024) apply event-study methods to an alternative measure of the in-person share and find support for a causal effect of school policy on (only) married mothers. In our context, the analogue to the pre-trends test—the placebo test—fails. Nevertheless, the failure lends a sharp interpretation to the OLS estimates and enables us to derive from them substantial insight into parental time use.

Consistent with this reading of results, our estimates from more standard specifications tend to exceed those in the broader literature on the effect 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 weaker responses of married mothers. The international evidence is more varied, but few if any find larger estimates than we report. Several studies find a comparable impact of longer school instruction for one parental group (i.e., married mothers) but not for others (Contreras and

Sepúlveda, 2017; Padilla-Romo and Cabrera Hernandez, 2019; Berthelon et al., 2023). Null effects have also been reported (Felfe et al., 2016).[2]

Next, our analysis of the ATUS contributes to a growing research agenda on telework. Pabilonia and Vernon (2023) document that take-up of telework increased at the onset of the pandemic, especially for mothers of children under the age of 13. Teleworking parents spent a large share of their day on secondary childcare activities. Atalay (2023) shows that these shifts were more pronounced for college-educated parents (see also Cowan, 2024). Our results echo these findings on the incidence of telework and caregiving during the pandemic. We extend this research by more precisely linking parental time use patterns to local in-person instruction shares.

Finally, we connect pandemic-era research on school closures to economic theory. We show analytically how our regression results inform models of parental investments and adolescent development and illustrate their broader implications for policy interventions and labor market dynamics. The mechanisms that we highlight—most notably, nonparental care—may in turn shed light on earlier empirical analysis of childcare availability (see above). In addition, we offer a means to formalize a new mechanism, telework, that is still used widely (see Barrero et al., 2024). We view our efforts to draw out lessons from the data within simple models as complementary to the estimation of richer models (see Del Boca et al., 2014, and Berlinski et al., 2024).

1. Data

This section introduces our measures of in-person instruction as well as our data sources for labor supply and other variables used in the regression analysis.

[2] 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., 2024). Our analysis will span all of 2020-21 and with more of a focus on the period beginning with the fall 2020 to spring 2021 school year.

1.1 In-person instruction

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 do vary across the trackers, which suggests that the different choices of methodology and sampling do 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, 2024; Hansen et al., 2024). 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.[3]

Our main measure of school policy from Parolin and Lee is constructed as follows. First, a school is classified as "closed" in some month m (and year $y \geq 2020$) if the number of visits to that school is down by at least 50 percent relative to month m in 2019. Parolin and Lee then calculate the closed share of schools within each county (and month). The complement of this—that is, one minus their figure—can be interpreted, roughly, as the in-person instruction share.

SafeGraph has several advantages. 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 classify the district according to one of a few coarse, discrete formats, such as "hybrid" or "virtual," whereas SafeGraph's data implicitly aggregate these modes into a single estimate of the change in on-site activity. Thus, SafeGraph provides unique breadth and precision.

The aggregation over foot traffic means, however, that SafeGraph captures both the *provision* of on-site instruction and parents' *take-up* of the in-person option.[4] The take-up decision is endogenous to labor supply: A parent who wants to work is more likely to enroll children in in-person instruction. For this reason, our SafeGraph-based estimates of the hours response to school closures should provide an upper bound. Estimates off CSDH data are subject to the same concern since the latter is derived from enrollment in each instruction mode. By contrast, Burbio documents

[3] Parolin and Lee drop private schools because their analysis uses other student data available only for public schools.
[4] 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.

only the availability of on-site instruction. Online Appendix C.2 shows that SafeGraph indeed yields the largest hours responses and Burbio the smallest; results from CSDH lie in between.

Geographic variation in in-person shares. Although Parolin and Lee's estimates cover the more than 3,000 U.S. counties, other data sources do not offer this same scope. The Current Population Survey, our source on hours worked, neither discloses school districts nor universally reports the respondent's county. Indeed, 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., 2024). First, we assign the county-level in-person share from Parolin and Lee to a survey respondent if the latter's county is one of the 280 identified in the CPS. Second, for respondents who have no county identifier but who belong to a disclosed MSA, we assign the mean in-person share among the non-identified counties in that MSA. Finally, we aggregate Parolin and Lee's estimates among those 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 whom 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 in-person shares. For each area, the average in-person instruction share in September-December 2020 is shown along the x-axis and the average share in January-May 2021 along the y-axis. The figure shows, first, that there are significant differences

[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. This step captures data from only 47 states because in 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

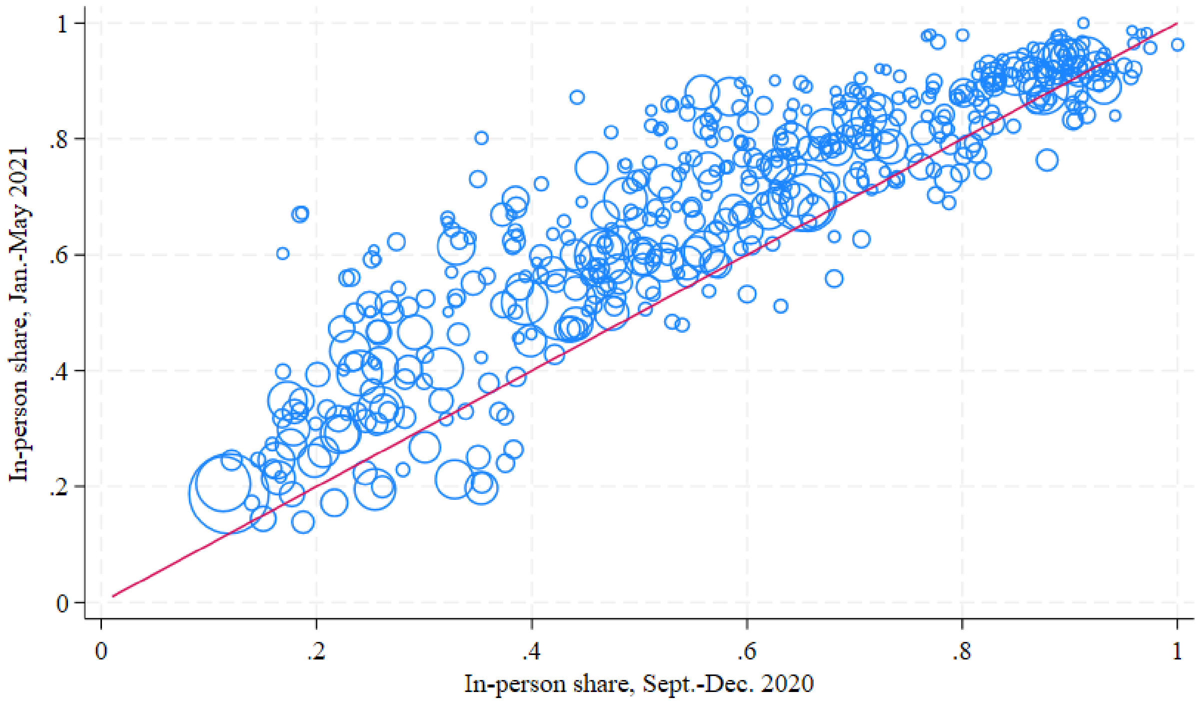

Note: This figure plots the average in-person share—from Parolin and Lee's (2021), measures derived from SafeGraph data—in September to December 2020 (x-axis) compared to January to May 2021 (y-axis). The size of each circle is proportional to the population in the geographic area.

across areas. In each of the two semesters, in-person shares span a wide range from 0.2 to 1. Second, these regional differences are, to some extent, persistent: in almost half of the areas, the in-person share shifted by less than 10 percentage points across semesters. In the other half of the areas, there was more substantial variation in instruction format *within* region. The latter variation generally reflected differences in the timing of reinstating in-person instruction in spring 2021.

What might account for the differences in in-person shares 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, both in-person instruction and labor supply might rise, even if the former has no causal effect on the latter.

In fact, the link between instruction format and COVID-19 case counts is remarkably modest. Online Appendix A lays out the evidence on this point. We suspect that monthly changes in case counts are weakly correlated with changes in that area's 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. Around the same time, Fairfax County (Virginia) announced that it would not reinstate on-site instruction until 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, unlikely to react sharply to changes in the state of the pandemic.

Instead, as Online Appendix B illustrates, 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. 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 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 (which may be zero) but also report results for employment status. 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

[6] For results on partisanship and union strength, see Grossmann 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.

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 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. The top panel collects tabulations for women, whereas the bottom panel refers to men. We also present results for the 280 CPS-reported counties (left-hand side) and the full sample of 478 local areas (right-hand side). Finally, for each sample of locations, the table reports on three groups; adults 21 and over; adults in the narrower range of ages 21-59; and parents of school-age children. (The ages of parents are unrestricted, but nearly all fall within the range 21-59.) As discussed later, our regression sample consists of all areas but restricts attention to ages 21-59. It is instructive to contrast our preferred sample to the alternative groups in Table 1.

A few patterns in the data are noteworthy, if not necessarily unexpected. First, the sample of all adults ages 21 and over has fewer kids in the home, is less racially and ethnically diverse, and works less than the other two groups. In other words, this subsample is observationally quite different from parents of school-age children. By contrast, the sample of adults ages 21-59 is very similar to parents (with marital status the obvious exception). Next, CPS counties are relatively urban, educated, and ethnically diverse and adopted in-person instruction less often. Thus, the use of all local areas captures a broader sample of parents and school policies. Finally, well-known differences in employment and marriage rates between mothers and fathers are evident in the table (Doepke and Tertilt, 2016). The labor supply of single mothers will be of special interest below.

---

[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. Accessed August 2, 2023.

Table 1: Summary Statistics

Variable	Women					
	CPS Counties			All Local Areas		
	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.	314,530	201,720	66,039	762,718	481,485	165,625
	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	0.643	0.649
Number of obs.	282,721	189,026	52,846	695,582	456,655	134,294

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 refers to the number of COVID-19 cases in the local area of the respondent in the survey month. In-person instruction refers to the share of schools in a local area open to in-person instruction in the survey month.

2. Empirical framework

Our aim is to examine the effect of in-person instruction on parental labor supply. We describe a series of specifications that differ only with respect to their treatment of unobserved heterogeneity.

We first consider the regression specification adopted in much of the related literature (Garcia and Cowan, 2024; Heggeness and Suri, 2021; Collins et al., 2021). This specification allows that the in-person share is endogenous to the state of the labor market but assumes it is (as good as) random with respect to parents' *relative* hours worked (that is, relative to the hours worked of childless adults). Thus, the regression leverages within-area differences in hours worked across adults with and without children.

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 having 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,

$$h_{iat} = \eta \mathbb{k}_{it} + \delta p_{at} + \psi p_{at} \mathbb{k}_{it} + \mathbf{\Gamma}' \mathbf{X}_{it} + \chi_a + \tau_t + \varepsilon_{iat}, \tag{1}$$

where h_{iat} is labor input of individual i in area a in month t. We generally take h_{iat} to be weekly hours worked, but we also present results where h_{iat} is a binary indicator of employment. The vector $\mathbf{X}_{it}$ captures additional individual-level controls to be described in the next section (and $\mathbf{\Gamma}$ is a conformable vector of coefficients); χ_a is an area fixed effect; and τ_t is a month fixed effect.[10]

[9] In practice, this share varies *within* area, 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.

The key parameter here is ψ, which measures the parental hours response to a unit difference in the in-person share.

Under certain conditions, ψ can be estimated consistently even if schooling mode is endogenous to local area trends. These trends will be picked up by p_{at} and reflected in δ, which measures the average response of *all* adults. An estimate of $\delta \neq 0$ will emerge if, for instance, schooling mode coincides with a general return to "normalcy", which shapes market-wide labor supply and demand. By contrast, ψ reflects the behavior of *parents' relative* hours of work (that is, relative to that of childless adults). Thus, the identifying assumption behind equation (1) is that any residual factors driving parents' relative hours are uncorrelated with local school closures.

Using a second specification, though, we can partially relax this identifying restriction. Consider the estimating equation,

$$h_{iat} = \delta p_{at} + \psi p_{at}\mathbb{k}_{it} + \zeta_a\mathbb{k}_{it} + \theta_t\mathbb{k}_{it} + \mathbf{\Gamma}'\mathbf{X}_{iat} + \chi_a + \tau_t + \omega_{iat}, \quad (2)$$

which introduces two new fixed effects that interact with parental status. (These replace and extend the regressor, $\eta\mathbb{k}_{it}$, in equation (1).) The parameter θ_t captures parent-specific factors behind hours worked that are common across areas but vary over time, whereas ζ_a captures fixed cross-area differences in parents' relative hours. These fixed effects allow that the (unobserved) labor supply motives of parents may evolve over time coincident with the "typical" school policy in the U.S. (i.e., θ_t) and/or correlate with the average policy in their local area (i.e., ζ_a). The identifying assumption underlying equation (2) is that changes in these idiosyncratic motives *within* a local area are uncorrelated with changes in *that area's* instruction format.

Thus, the added controls in equation (2) narrow the scope of variation used to identify the coefficient, ψ. Equation (2) recovers a significant effect only to the extent that parents' relative hours co-move with the in-person share in their area.[11] By contrast, equation (1) exploits both the within- and across-area correlation of in-person shares and parents' relative hours. Equation (2) offers potentially more credible identification but at the cost of statistical power.

Extending the approach underlying equation (2) still further, one could insert *individual* level fixed effects. This specification maps the change in in-person shares facing each survey respondent to the change in her hours of work, thereby identifying ψ using only variation in school policies over time. An individual fixed effects regression can be estimated by using the longitudinal dimension of our data. Since respondents in the same area are exposed to the same school policies, though, this specification yields results that are similar to what is implied by the introduction of parent-by-area controls in equation (2).

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 of equations (1) and (2) 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 considered a parent of a school-age child if they have a child between ages 5 and 17 in the home. Households with only children 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. Online Appendix C.3 considers several variations on this sample. First, we divide households more finely by age of the eldest child and show that our results below largely stem from households with children under age 13. Second, we find that the inclusion of childless adults over age 59 yields

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

larger parental labor supply responses. However, the parental hours response in this context reflects—and is amplified 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.5. 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 (2024). 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 in an industry and occupation. Unfortunately, these data are not reported in the CPS for most labor force

[12] The only controls here that are not present in Garcia and Cowan (2024) are the indicators for rural-urban-suburban status and for the presence of under-five-year-old children in the home.

nonparticipants.[13] Nevertheless, Online Appendix C.6 does introduce these controls and confirms that the impact of in-person shares is estimated to be even smaller than reported below.[14]

3.1 Full sample

We proceed to estimate the standard two-way fixed effects model in equation (1), with weekly hours worked as the dependent variable. Online Appendix C.1 reports results for employment. Table 2 presents results for two periods: the longer one spans all of 2020-21 except for the summer months, whereas the shorter period covers the 2020-21 school year (September 2020 – May 2021). Note that the former period featured school closures, whereas the latter was characterized by a staggered reopening to in-person learning. Thus, by separating out the latter period, we can examine if closing and reopening had meaningfully different effects. For each period and each outcome, we also report results separately for men and women. Finally, in view of the arguments in Solon et al. (2015), we report both unweighted estimates and estimates that apply CPS sample weights. While weighting makes little difference on balance, we will highlight the few instances where it does.

Consider first the results for the longer sample period (2020-21). The main parameter of interest is ψ, which measures the response of parents' hours worked to the in-person share. Among women, a shift from fully virtual to fully in-person instruction implies an increase in hours worked of roughly 0.5 per week. The overall hours response among fathers is nearly identical in the unweighted specification, but closer to 0.2 when applying the CPS sample weights. Finally, estimates of δ, which capture the area-wide hours response, are positive and significant, suggesting

[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. The ORGs *as a whole* make up only one quarter of the CPS sample.

[14] Relatedly, we also do not restrict the sample based on industry or occupation affiliation. There is arguably a case to exclude respondents in the education sector since changes in in-person shares might mechanically imply changes in their hours worked (see Hansen et al., 2024). We confirm that this restriction has a negligible impact on our estimates.

Table 2: Estimates of Equation (1)

	All of 2020-21		2020-21 School Year	
	Women:			
In-person share, δ	1.205***	1.130***	-0.853	-0.931
	[0.338]	[0.401]	[0.594]	[0.690]
In-person × kids, ψ	0.582*	0.472	2.113***	2.293***
	[0.304]	[0.326]	[0.590]	[0.629]
Number of obs.	447,899	447,277	228,550	228,225
	Men: Weekly Hours			
In-person share, δ	1.285***	0.856**	-0.009	0.006
	[0.382]	[0.430]	[0.647]	[0.696]
In-person × kids, ψ	0.566*	0.210	1.456**	1.315**
	[0.315]	[0.321]	[0.589]	[0.559]
Number of obs.	432,856	428,244	221,080	218,575
CPS Weights	No	Yes	No	Yes

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.

that in-person shares may pick up broader shifts in the propensity to work.[15]

Next, we turn to the results in Table 2 for the 2020-21 school year. These results paint 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. Finally, estimates of δ are no longer significant. We have confirmed that these differences across the two periods reflect the influence of the months that preceded the 2020-21 school year (January – May 2020) and *not* the months that followed (September – December 2021).

[15] Appendix C.1 shows that the estimated changes in hours worked reflects an extensive margin adjustment for mothers and an intensive margin adjustment for fathers.

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 initial area-wide reactions to these disruptions will be reflected in a significant response to (correlated) changes in in-person shares. Thus, the absence of controls for such trends may lead to different estimates of ψ across different periods.

Table 3: Estimates of Equation (2)

	Women: All of 2020-2021			
In-person × kids, ψ	2.359***	2.501***	-0.051	0.096
	[0.634]	[0.654]	[0.672]	[0.751]
Number of obs.	447,899	447,277	447,899	447,277
	Women: 2020-21 School Year			
In-person × kids, ψ	2.458***	2.568***	-0.131	-0.440
	[0.633]	[0.668]	[1.127]	[1.248]
Number of obs.	228,550	228,225	228,550	228,225
	Men: All of 2020-21			
In-person × kids, ψ	1.886***	1.708***	-0.051	-0.239
	[0.645]	[0.602]	[0.705]	[0.812]
Number of obs.	432,856	428,244	432,856	428,244
	Men: 2020-21 School Year			
In-person × kids, ψ	1.778***	1.596***	-1.695	-1.129
	[0.629]	[0.590]	[1.191]	[1.344]
Number of obs.	221,080	218,575	221,080	218,575
CPS Weights	No	Yes	No	Yes
Month × parent F.E.	Yes	Yes	Yes	Yes
Area × parent F.E.	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. Standard errors are clustered at the geographic area level. "All 20-21" pools data for all of 2020 and 2021 exclusive of 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.

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

In view of this concern, we re-run the regression with additional controls for parent-specific trends in labor supply. Formally, these trends are modeled as parental status-by-month fixed effects ($\theta_t \mathbb{k}_{it}$ in equation (2)). The first two columns of Table 3 report the results. (With parent-by-month effects, estimates of δ are now insignificant in the full sample and are omitted here.) Under this specification, the adjustment of parents' hours to in-person instruction is now remarkably stable across time. Among mothers, a shift from fully virtual to fully in-person instruction yields an increase in weekly hours of around 2.4 to 2.6—regardless of the sample period. The response among fathers is somewhat smaller—weekly hours increase by around 1.6 to 1.9, depending on the weighting—but again, is virtually unchanged across sample periods. Thus, as anticipated, the parameter instability in Table 2 reflects 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. [17]

Just as there may be parent-specific trends in hours worked, there may be parent-specific factors behind average hours in a given area. These factors drive a wedge between the mean hours of parents and childless adults within an area and may vary across areas. Such spatial differences pose a threat to identification if they are correlated with (average) 2020-21 in-person instruction rates. The reasons for such a correlation are perhaps not immediate (we return to this shortly), but it is easy all the same to add controls for spatial heterogeneity. As previewed in Section 2, these controls take the form of parental status-by-area fixed effects ($\zeta_a \mathbb{k}_{it}$ in equation (2)).

The impact of these controls, shown in the final two columns of Table 3, 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

[17] Online Appendix C.1 confirms that, in regressions with parent-by-month effects, the extensive margin continues to play an outsized role in women's labor supply response but matters little for men.

principally off cross-area comparisons of parents' relative hours worked. With additional controls for average regional differences in labor supply, estimated effects of the in-person share disappear.

Online Appendix C.7 shows that the introduction of individual fixed effects has a similar impact as the parent-by-area regressors. Intuitively, each set of controls isolates variation in in-person shares over time within a fixed unit (either a person or area). This variation alone does not identify a statistically significant effect of in-person instruction.

One could question, though, if we have "over-controlled" for unobserved heterogeneity. Even if regional differences in average school policy were exogenous, the addition of the parental status-by-area terms alone could capture much of this variation. To assess the need for these controls, consider a simple placebo test. Suppose in-person shares in 2020-21 are correlated with long-run regional differences in relative parental hours. It follows that average policies in the pandemic should predict *pre*-pandemic labor supply.

In fact, this "pre-trend"—the correlation between the pandemic-era instruction format and pre-pandemic hours—is evident in the raw data. Figure 2 illustrates this point. The x-axis shows the average in-person share in each of our local labor market areas over 2020-21. The y-axis is based on pre-pandemic hours data from the CPS. Specifically, it shows the local-area average of parents' hours less average hours of childless adults over the five years before the pandemic, 2015-19. The left panel reports results for mothers, and the right panel pertains to fathers. Remarkably, parents' relative labor supply in the pre-pandemic period appears to be several hours higher in areas where instruction was largely in-person in 2020-21 than in areas where it was largely remote. To pursue this point further, we apply equation (1) to test if in-person shares in 2020-21 predict pre-pandemic hours. The sample is drawn from the CPS and consists of adults ages 21-59 in the years 2015-19. All individual-level control variables described above are included. The schooling mode, which was formerly measured by monthly data on in-person shares in 2020-21 (p_{at}), is now

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

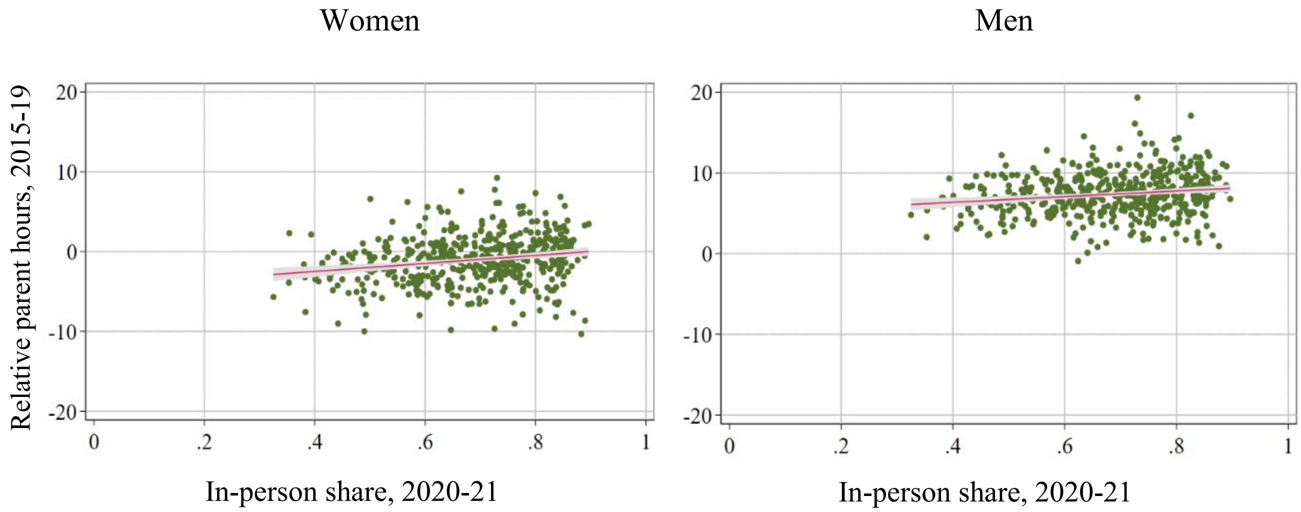

Note: This figure plots (on the y-axis) the difference in average pre-pandemic weekly hours between parents and childless adults against (on the x-axis) the average in-person share in the pandemic period. Each marker is a local labor market area. The left panel is based on hours data among women ages 21-59; the right panel refers to 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 5.024 (s.e. of 1.106), and the line of best fit in the right panel (among men) has slope 3.504 (s.e. of 0.976). To mitigate sampling error, we drop the seven areas with fewer than 50 mothers or fewer than 50 childless women (left panel) and the nine areas with fewer than 50 fathers or fewer than 50 childless men (right panel).

the area-level *mean* of the latter and denoted by p_a. The regressor of interest is the interaction term, $p_a \mathbb{k}_{it}$. (We do not include p_a as a stand-alone regressor, since it is absorbed by area fixed effects.) A significant coefficient on the interaction means that average on-site shares in 2020-21 pick up general regional differences in parents' relative hours in 2015-19.[18]

The regression analysis corroborates that the pandemic-era schooling mode is strongly related to pre-pandemic *maternal* hours but uncovers a weaker connection to paternal labor supply. These estimates are detailed in Online Appendix B and summarized here. The results are most striking when using an area-level mean p_a based on all of 2020-21 (excluding summer months). In areas that selected full-time in-person instruction, mothers' relative labor input *prior* to the pandemic is estimated to be roughly 3.4 weekly hours higher than in areas with full-time virtual

[18] Results are virtually unaffected if we insert parental status-by-month effects as in equation (2).

instruction. Among fathers, in-person instruction implied around one more hour of work per week, although the latter is not statistically significant. 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). Alternatively, if we compute mean in-person shares based on 2020-21 school year data, the estimate for mothers falls to about two hours per week but remains strongly significant. The analogue for men lies between 0.6 and 0.8 hours per week but is, again, not significant. In the Appendix, we find the same pattern of results with in-person share measures other than SafeGraph.

To reflect on these results, 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 the model in Section 5.) The first is the cost of school-age childcare. The second is commute time to work, which reduces, all else equal, time spent with children.

Online Appendix B shows that commute times and school-age childcare costs are correlated with (pandemic-era) in-person shares. This connection runs, in part, through their association with local partisan affiliation. As we noted, in-person shares were highest in areas that heavily supported Donald Trump. At the same time, commutes are longer in metro areas where Trump's vote share was low. Higher childcare costs in anti-Trump areas may partly reflect the burden of higher minimum staff-to-child ratios, suggesting a greater propensity to regulate.

In addition, the Appendix reports on the connection between commute length and childcare costs, on the one hand, and parental labor supply on the other. A statistically significant correlation suggests that any *other* outcome related to commute times and childcare prices, such as the in-person share, will emerge as an apparent contributor to parents' labor supply. We find that longer commutes and higher childcare prices are indeed associated with lower maternal hours worked. By contrast, paternal hours are less sensitive to local childcare costs and essentially uncorrelated with commute times. In qualitative terms, these results echo more careful, causal analyses (see Black et al. (2014) on commute times and Mumford et al. (2020) on childcare prices).[19]

Taking stock of our findings, we conclude with the following observations. The results of the placebo test demonstrate that equation (1) fails to address the endogeneity of schooling mode. As a result, equation (1) likely yields an upwardly biased estimate of its effect. However, the source of this endogeneity is not fully resolved. Among mothers, the connection between hours worked, commute times, childcare prices, and in-person shares suggests that schooling mode stands in for more fundamental forces in the local area. This narrative does not apply neatly to fathers, though.

3.2 Education and marital status

In line with related research, we next ask if parental labor supply responses to virtual instruction 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. The regression model retains parental status-by-time effects but excludes parental status-by-area effects. We confirm that the inclusion of the latter eliminates the statistical significance of the estimates, just as they do in Section 3.1. One might then view the results below as the strongest case that one could present for a role of schooling mode in parental labor supply.

Education. We first divide our sample into a noncollege group—workers with less than a four-year degree—and workers who completed college. We then further split each of these two groups by gender. Results are reported in Table 4.

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

[19] For a review of research on childcare prices and maternal labor supply, see Blau and Currie (2006). The divergence between paternal and maternal labor supply responses is a feature of the lifecycle model in Guner et al. (2020).

Table 4: Estimates by Educational Background

	Noncollege		College	
	Women			
In-person × kids, ψ	2.074***	2.374***	1.818*	1.851*
	[0.771]	[0.849]	[1.001]	[1.082]
Number of obs.	266,258	265,968	181,641	181,309
	Men			
In-person × kids, ψ	1.999***	1.790***	1.078	1.152
	[0.751]	[0.666]	[0.863]	[0.973]
Number of obs.	284,723	281,867	148,133	146,377
CPS Weights	No	Yes	No	Yes

Note: Each column within each panel is a separate regression, and the column header reports the regression sample (i.e., "noncollege women"). The period is all of 2020-21 but with summer months excluded. A college (noncollege) graduate is one who did (not) complete a four-year degree. 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.

not statistically distinguishable from one another. Thus, among women, college experience is not a strong predictor of the labor supply response to the in-person share.

The education gradient among men is somewhat more evident. The college educated do not significantly adjust their hours in response to variation in the in-person share. By contrast, the response of noncollege men is similar to that 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, as shown in Online Appendix C.8. Mothers in these households raise labor supply by up to one hour more than shown in Table 4, whereas fathers' behavior is in line with the college group as a whole. This imbalance between spouses is evident only in college-educated couples. In households with noncollege-educated parents, spouses' hours responses are almost identical.[20]

[20] These patterns do not seem to reflect intra-household differences in earnings opportunities: a college-educated father is *no* more likely than a noncollege graduate to have higher earnings than his spouse. See Online Appendix C.8.

Table 5: Estimates by Marital Status

	Married		Unmarried		Lone adults	
	Women					
In-person × kids, ψ	2.256***	2.902***	2.591**	2.047**	3.995***	2.973**
	[0.788]	[0.836]	[1.032]	[0.948]	[1.296]	[1.324]
Number of obs.	242,743	242,351	205,156	204,926	60,291	60,282
	Men					
In-person × kids, ψ	1.824***	2.256***	1.657	1.254	4.072*	3.033
	[0.661]	[0.709]	[1.455]	[1.540]	[2.266]	[2.498]
Number of obs.	223,471	219,663	209,385	208,581	55,284	55,275
CPS Weights	No	Yes	No	Yes	No	Yes

Note: Each column within each panel is a separate regression, and the column header reports the regression sample (i.e., "married women"). The period is all of 2020-21 but with summer months excluded. A "lone adult" is a respondent who does not live with any other individual age 18 or over. *** 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.

Marital status. We next split the sample by marital status. In addition, within the unmarried, we look at households where the parent is the lone adult. The labor supply response of a single parent is likely to depend on the household's composition. For instance, a parent in a coresidential arrangement with other adults may receive steadier childcare support than a lone-adult parent. This consideration is empirically relevant: almost 60 percent of unmarried mothers live with at least one other adult, which includes unmarried partners, parents, and older children.

Our estimates in Table 5 confirm that household composition mediates the role of marital status. The response of hours worked among all unmarried mothers ranges from 2.0 to 2.6, which is not too different from that of the married sample. However, this estimate masks the difference between mothers with and without other adults in the household. Among lone-adult mothers, hours worked are more responsive: a shift from a virtual to in-person format implies an increase of 3.0 to 4.0 weekly hours. By contrast, the response of unmarried women in co-residential arrangements (not shown) is 1.5 hours and statistically insignificant. The narrative for men is broadly similar although the estimates are less precise (in part because few unmarried men live with their children).

It is instructive to compare results in Table 5 with other research in this area. Garcia and Cowan (2024) adopt a specification very much like equation (1). Our estimates are comparable to, or higher than, theirs save for unmarried men (for whom our ψ is one hour lower). One distinguishing feature of our specification is the use of parental status-by-month effects, which tends to elevate estimates of ψ in the full sample 2020-21. Other differences between our specifications offset one another to some degree. Specifically, the presence of older respondents in Garcia and Cowan's sample elevates their estimates (see Online Appendix C.4), but their restriction to CPS counties and inclusion of industry and occupation controls reduces them (see Online Appendices C.5 and C.6).[21] Hansen et al. (2024) also consider a specification akin to equation (1) but show that their results are robust to event-study methods that abstract from the spatial variation that underlies our placebo test. They uncover a statistically significant effect for married mothers but not for unmarried mothers or married fathers.[22] Thus, we generally find larger estimates of the labor supply response and yet we will argue in Section 5 that even our results are 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) (Flood et al., 2023). Our ATUS sample is selected to conform to the extent possible with

[21] A more complete mapping from Gracia and Cowan's results to our own, including the effects of each of these specification choices, was included in an earlier version of this paper and is available upon request.
[22] The point estimates in Hansen et al. (2024) are not quite comparable to ours because the authors develop their own in-person shares based on SafeGraph. We and Garcia and Cowan use those from Parolin and Lee (2021). In our data, our placebo test fails for women in all demographic groups.

our treatment of the CPS. Therefore, we again restrict attention to individuals ages 21-59 who are childless adults or parents of school-age children. The sample period covers 2020-21 but for the period mid-March to mid-May 2020 during which field work was suspended due to the pandemic.

For each respondent, we observe a minute-by-minute diary of a single day that describes how, where, and with whom they spent their time. However, the days of the week are not uniformly represented: Half come from Saturday or Sunday. We implement a simple reweighting that mimics a uniform sample over days of the week.[23] Alternatively, the oversample of weekend days can be corrected by use of ATUS sample weights. We present results based on both weighting schemes.

Our analysis addresses time allocation across several dimensions. Each respondent's diary entry is assigned a detailed activity code, and we group activities into a few broad categories: work, leisure, home production, childcare, commute time, and sleep. We then estimate how hours spent in each category respond 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. Finally, 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.

Remarkably, the reinstatement of in-person instruction has, on the whole, no significant impact on any major category of time use, from work to leisure and home production. These results, which are reported in Online Appendix D, hold for the full sample and when we split the data by college attainment. Given the modest size of our sample, what we take from this exercise is that, whatever are the "true" effects of schooling mode on time allocation, they are not large enough to detect in the ATUS.[24]

[23] To illustrate, if Saturday represents 1/4 of the sample, we apply a weight of 4/7 to all Saturday observations.
[24] The discrepancy between CPS and ATUS results on hours worked is unlikely to reflect systematically different measurements. Research has found substantial agreement between the two sources (Frazis and Stewart, 2004, 2014).

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 sheds light on how an activity was performed. For instance, while *total* market time might be unresponsive to the closure of in-person instruction, a greater *share* of it may overlap with childcare. Fortunately, the ATUS asks if there was a child in the respondent's care, even if the respondent was engaged in another activity. (We refer to this childcare time as a secondary activity.)[25] Thus, we can observe if parents supervise school-age children while they work at home.

Table 6 reports on the role of working from home as a means of supplying both childcare and market time. To start, the first two columns reiterate that total working hours in the ATUS are insensitive to instruction format. The next two columns report results for *total* hours working at home. Interestingly, this, too, does not respond significantly. However, the fifth and sixth columns show that time spent working at home (as the primary activity) while caring for children (as the secondary activity) 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 6-7 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 extent from college educated mothers, but standard errors in these subsamples are rather large (which is why we pool men and women in Table 6). 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).

[25] Note that childcare is the only activity that can be recorded as a secondary activity. The ATUS does not ask survey respondents for secondary activities outside of childcare.

In the final columns of Table 6, we report how school closures alter the *total* time spent with children, which includes both primary and secondary childcare.[26] Overall, local school closures led to an increase of 14 to 19 hours per week with one's children, with much of the latter due to the response of the college-educated. This estimate captures adjustments on the work-from-home margin (columns five and six) as well as variation in the extent to which childcare overlapped with *non*-market activities (e.g., leisure and home production).

These results strongly suggest that college educated parents relied, in part, on telework to sustain their hours worked when instruction was virtual. Nevertheless, we would not necessarily infer from Table 6 that the labor supply of the college educated would have fallen 6 weekly hours *but for* telework. The reason is that the noncollege educated in the ATUS also smoothed hours worked but did not rely on telework. Thus, in the absence of telework, the college group would have presumably taken up, at least to some degree, measures adopted by the noncollege group to cope with shifts in instruction format.[27]

We next turn to one of these other possible margins of adjustment: the utilization of nonparental childcare. A survey fielded in late 2020 by Calarco et al. (2021), and analyzed further in Yang et al. (2025), 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). Sixty percent of surveyed families reported using informal care, which included help supervising children learning at home in fall 2020. By excluding spring 2020, though, the survey likely does miss the disruptions faced by many caregiving arrangements at that time. Even during those initial

[26] These results are not strictly comparable to several others in the table. The reason is that the measured outcomes such as "work" and "work at home" do not capture the time spent in those tasks as secondary activities.
[27] To recover causal effects more credibly, one could try to exploit plausibly exogenous variation in workers' access to telework. However, measures of access are based on occupations (Dingel and Neiman, 2020) and are not easily mapped to nonemployed survey respondents.

Table 6: Work at Home, Childcare, and Instruction Format

	Work		Work at Home		Work at Home + Childcare as Secondary Activity		Childcare, Primary or Secondary Activity	
All								
In-person × kids, ψ	-0.519	-5.182	-3.124	-2.852	-5.937***	-4.796***	-18.912***	-14.003**
	[4.091]	[5.223]	[3.961]	[4.756]	[1.466]	[1.626]	[4.903]	[5.582]
Non-College								
In-person × kids, ψ	-2.791	-7.840	1.528	1.192	-1.521	-1.009	-9.683	-2.211
	[7.040]	[7.921]	[4.181]	[4.741]	[2.154]	[2.024]	[8.374]	[9.237]
College								
In-person × kids, ψ	1.736	0.467	-3.058	-3.901	-7.328***	-6.432**	-25.149***	-25.572***
	[5.867]	[6.024]	[6.892]	[7.875]	[2.620]	[2.801]	[5.053]	[5.265]
ATUS Weights	No	Yes	No	Yes	No	Yes	No	Yes

Note: Each column within each panel is a separate regression. The dependent variable is the implied number of hours per week spent in each activity. The panel title reports the regression sample. There are 6,622 observations in the first panel, 3,371 observations in the second panel, and 3,178 observations in the third panel. Relative to equation (1), we also include fixed effects for days of the week as well as parental status×month controls. 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.

Table 7: Time with Others' Children and Local School Formats

	All		Men		Women	
			All			
In-person share, δ	-1.832	-1.753	-2.849	-0.477	-1.597	-2.961
	[1.690]	[2.082]	[1.897]	[2.695]	[2.348]	[2.535]
Number of obs.	4,848	4,848	1,983	1,983	2,787	2,787
			Non-College			
In-person share, δ	-2.177	-4.261*	-0.083	-0.401	-3.249	-6.445**
	[2.157]	[2.287]	[2.837]	[2.626]	[2.790]	[3.105]
Number of obs.	2,945	2,945	1,106	1,106	1,725	1,725
			College			
In-person share, δ	-1.533	1.029	-6.476	-0.681	-0.041	1.957
	[3.019]	[3.734]	[5.448]	[6.511]	[4.396]	[4.491]
Number of obs.	1,817	1,817	765	765	952	952
ATUS Weights	No	Yes	No	Yes	No	Yes

Note: Each column within each panel a separate regression. The dependent variable is the implied number of hours per week spent with other's children. 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.

months of the pandemic, though, caregiving hours appear to have risen in households where older family members resided (Truskinovsky et al., 2022).

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 spent with children under age 18 who are *not* the respondent's son or daughter. This estimate excludes time spent at work in order to identify unpaid, informal care of the sort that a grandparent or other older relative might provide.

Table 7 reports how these hours of care vary with the in-person share of instruction. Note that since the sample consists of only potential nonparental caregivers, the covariate of interest now is just the in-person share rather than the interaction of the latter with parental status. The identification assumption in this context is that schooling mode did not systematically vary with older respondents' preferences or opportunities for caregiving. Estimates from the ATUS suggest that older respondents' caregiving was responsive to the in-person share. In the full sample, the resumption of in-person instruction implies a reduction of nearly 1.8 hours per week in the time older respondents spend with children, though this estimate is not statistically significant. We obtain larger estimates if we consider those without a college degree: their weekly hours of caregiving fall by up to 4.3 when on-site instruction returns. The response among noncollege women appears to be even larger. One way to interpret these results is to view the grandparent's college experience as a proxy for that of the parent, which suggests that noncollege households relied more on nonparental care.[28] This interpretation is consistent with Kwon (2024), who finds higher parental hours in CPS households where grandparents were present. Moreover, her estimates are largest for households with lesser educated parents.

5. Discussion

We now use a series of time allocation models to guide a discussion of our regression results. We first consider a very simple set-up where a single parent faces a one-for-one tradeoff (in time) between labor supply and childcare. Under a reasonable parameterization, the model implies labor supply responses that far exceed any reported estimate. We then illustrate 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 the noncollege educated, who were less likely to access telework. 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.

A simple baseline. A single parent maximizes utility over consumption, leisure, and child development subject to two constraints on her time. The first constraint is that the allocation

[28] On the intergenerational correlation of educational attainment, see Kane (1994) and Cameron and Heckman (2001).

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 is 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 time, but it arguably captures the dimension of in-person instruction that is most relevant to parental labor supply. We assume that a child who is not in school must be under the parent's supervision. We introduce private nonparental care below.

Formally, the time constraints are as follows. 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,

$$l + m + n = 1, \tag{3}$$

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

$$g + m = 1. \tag{4}$$

Together, equations (3) and (4) imply $l = g - n$: a decrease in on-site instruction time, g, lowers leisure one for one unless market hours fall.

We assume the parent takes g as given. This rules out substitution from an institution with only virtual instruction to one that is in person. Where this did occur in practice, it appears to have involved a switch from public to private school.[29] For the typical parent, though, the cost of such a switch was likely prohibitive. Therefore, we focus here on other margins of adjustment.

We assume that period utility is given by

$$u(c, l, q) = \alpha \ln c + \beta \ln l + (1 - \alpha - \beta) \ln q, \tag{6}$$

[29] However, much of the 3 percent decline in public school enrollment in Fall 2020 reflected increased homeschooling (Musaddiq et al., 2022; Bacher-Hicks et al., 2024). Our model interprets this as more time under parental care.

where $\alpha, \beta \in (0,1)$. The Cobb-Douglas specification follows Berlinski et al. (2024) and is the form of period utility often used in models of home production. In our context, 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. Since $m = 1 - g$, though, $q(g, m) = q(g, 1 - g)$: q is pinned down by g, which is taken as given. A more substantive choice problem for q will emerge when we introduce another source of supervision: a form of private nonparental care. Nevertheless, the level of g still shapes labor supply, n, via the time constraint (5).

Time allocations are divisible and, therefore, the model will yield only interior solutions. In practice, though, the hours responses of parents also reflect movements on the extensive margin. In our view, what we sacrifice in realism is worth the insight that it affords. The comparative statics with respect to local changes in g can help reveal fundamental economic forces at play (even if observed shifts in market and on-site time tend to be "lumpy").[30]

Initial comparative statics. The first-order condition for leisure implies

$$l = g - n = \frac{\beta}{\lambda w}, \tag{7}$$

where λ is the marginal utility of consumption and w is the wage. Suppose for now that households can insure consumption to the extent that λ is invariant to g. It follows from equation (7) that market hours move one-for-one with on-site time. Intuitively, the demand for leisure does not change since its price is pinned down by λ and w. Therefore, n must fully offset a shift in g.

The assumption of perfect insurance is of course somewhat stylized, although the surge of government transfers likely did enable households to smooth consumption to a considerable extent

[30] Key features of the model, such as the curvature over leisure, are also likely to bear on the extensive margin. For instance, suppose a worker chooses $n = 0$ or $n = N > 0$ and derives utility $\beta v(g - n)$ where v is concave and β is heterogeneous. The employment *rate* varies inversely with the value of foregone leisure, $v(g) - v(g - N)$ (Mulligan, 2001). A lower g raises this value, and reduces labor supply, to an extent that depends on the curvature of v.

(Wu et al., 2022). Nevertheless, as an alternative, suppose parents live "hand to mouth." Therefore, consumption must satisfy $c = wn$. It follows that $\lambda = \alpha/c = \alpha/wn$, and equation (7) becomes

$$l = g - n = \frac{\beta}{\alpha} n. \tag{8}$$

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

$$\frac{\mathrm{d}n}{\mathrm{d}g} = \frac{1}{1 + \beta/\alpha}. \tag{9}$$

Equation (8) says that β/α is identified by the ratio of leisure to market work time, which can be calculated from data in the American Time Use Survey. We report two figures that bridge different approaches to the measurement of leisure (see Aguiar et al., 2012). First, if all sleep is excluded from leisure, we find that $\beta/\alpha = 1.1$, which implies that an hour more of in-person instruction yields approximately 0.5 more hours of market work. Alternatively, we treat sleep time beyond 6 hours as leisure. This approach elevates β/α and yields $\mathrm{d}n/\mathrm{d}g \approx 1/3$.

This prediction (far) exceeds estimates reported here or elsewhere in the literature. With the reinstatement of 33 hours of on-site instruction (U.S. Department of Education, 2008), equation (9) predicts that a shift from virtual to in-person will lift labor supply by 16 hours per week. By contrast, our OLS estimates suggest a labor supply response between 2-4 hours per week, i.e., at most $\mathrm{d}n/\mathrm{d}g \approx 0.1$ In this sense, our regression estimates are unexpectedly small.[31]

Telework. A key assumption embedded in equation (3) is that parents cannot simultaneously perform market work while they supervise children. However, ATUS data suggest that telework enabled (at least college-educated) parents to provide some childcare even as they continued to work. We illustrate a tractable way to capture this notion of telework in the model.

[31] Alternatively, if some parents will not work in any state of the world, the average "treatment" amounts to an increase in on-site time less than 33 hours. Suppose we discount 33 by 25%, which matches the mean nonemployment rate in Table 1. Still, given a labor supply response of around 3 hours, $\mathrm{d}n/\mathrm{d}g = 3/\big(33 \times (1 - 0.25)\big) = 0.12$.

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,

$$t(m, n) = (m^{\rho} + n^{\rho})^{1/\rho}, \tag{10}$$

where $\rho \geq 1$. The time constraint (3) 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). One might also want to allow leisure and childcare time to overlap, consistent with estimates in Table 6. We leave this for future research and focus here on the role of telework.

Equation (10) encompasses two polar cases. The first is $\rho = 1$, which recovers the original time constraint (3), $l + m + n = 1$. 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 (10) will be important. First, equation (10) 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 (i.e., t is convex) and decreases in the time allocated to the other activity (i.e., $\partial^2 t / \partial n \partial m = \partial^2 t / \partial m \partial n < 0$). The intuition is that, if m is large relative to n, a parent can identify many childcare tasks that can be done concurrently with more market work but few work tasks that can be done jointly with more childcare. Therefore, the time price of another hour of work is small, but the price of another hour of care is high.

These properties formalize the sense in which equation (10) 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 $t_n / t_m = (m/n)^{-(\rho-1)}$. Thus, at higher values of ρ, a one percent increase in childcare time (all else equal) yields a steeper decline in the relative price of market work.

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

$$l = 1 - t(m, n) = \frac{\beta}{\lambda w} \cdot \frac{\partial t}{\partial n}. \tag{11}$$

A decline in on-site time, g, now has two effects. The first is familiar: since parental time must rise, leisure would fall all else equal. To stem the decline in leisure, labor supply is reduced.[33] The second effect is novel: an increase in m also reduces the time price of market work, $\partial t / \partial n$. This stimulates *more* labor input, mitigating the decline in labor supply due to the former effect.

More formally, under perfect insurance ($\mathrm{d}\lambda = 0$), the comparative static is,

---

[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, to complete a new childcare task jointly with current market work. Therefore, $t_m = 1$.

[33] The *extent* to which it is reduced will depend on the shape of t. Thus, even the quantitative effect of this familiar mechanism is different under $\rho > 1$.

$$\frac{\mathrm{d}n}{\mathrm{d}g} = \frac{1-(\rho-1)/\phi(l)}{(n/m)^{\rho}+(\rho-1)/\phi(l)} \cdot \frac{n}{m}, \tag{12}$$

where $\phi(l) \equiv (1-l)/l$ and $m = 1-g$. When $\rho = 1$, equation (12) 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. In fact, there is a unique value of ρ, given by $\rho = 1 + \phi(l)$, that induces *no* change in market time. The term $\phi(l)$ captures the degree of curvature over l in the utility function: if $\phi(l)$ is large, (log) marginal utility of leisure rises steeply with any reduction in l, which prompts the parent to reduce market hours more substantially. For $\mathrm{d}n = 0$, the motive to multi-task, as parameterized by ρ, must be strong enough to match the force of this curvature.

To illustrate the implications of this result, consider the college educated, who relied on telework to sustain labor supply. From the ATUS, leisure for this group constitutes 38 percent of total time allocated to market work, childcare, and leisure.[34] Therefore, the observation $\mathrm{d}n \approx 0$ requires $\rho \approx 2.63$. More generally, we can identify conditions such that $\mathrm{d}n/\mathrm{d}g$ decreases in ρ, which provides a means to match an array of market hours outcomes. See Online Appendix E for a complete characterization.

Nonparental care. Thus far, we have assumed that a child must be supervised by her school or parent. However, changes in labor supply—and along other dimensions of time use—are relatively modest even among workers with little access to telework (i.e., the noncollege educated). One explanation for this is that parents turned to private nonparental care. Note that to zero in on this issue, we will abstract from telework in what follows. Online Appendix E shows that our main insights can be derived in a model that integrates both margins of adjustment.

[34] This is the notion of l within the model. Therefore, we abstract from other margins of time use for this calculation.

The introduction of nonparental care implies a simple, but potentially substantive, change in labor supply. If we denote time in private nonparental care by x, the analogue to equation (4) is

$$g + m + x = 1, \tag{13}$$

which says that a child is supervised by a school, parent, or private third party. The first-order condition for hours worked extends equation (7) to incorporate nonparental care,

$$n = g + x - \frac{\beta}{\lambda w}. \tag{14}$$

Market work now 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

$$\begin{gathered} q = g^{\gamma} \mathcal{q}(m, x)^{1-\gamma}, \text{ with} \\ \mathcal{q}(m, x) = (\mu^{1-\varphi} m^{\varphi} + (1-\mu)^{1-\varphi} x^{\varphi})^{1/\varphi} \end{gathered} \tag{15}$$

and where $\gamma \in (0,1)$ and $\varphi \leq 1$. Equation (15) uses a Cobb-Douglas outer nest to aggregate on-site instruction time (g) and a "bundle" of private care ($\mathcal{q}$). The latter inner bundle is a CES aggregate of parental (m) and private nonparental (x) care time. The parameter φ controls the elasticity of substitution between m and x, which is given by $(1-\varphi)^{-1}$. The CES form is a popular specification of development production functions (see, e.g., Cunha et al., 2010) and has been applied in the context of parental and nonparental private care (Berlinski et al., 2024).[35] The literature offers less guidance on the role of g in q. We opt for a Cobb-Douglas outer nest because

[35] Del Boca et al. (2014) use a Cobb-Douglas aggregator over all inputs but omit on-site time.

it simplifies the analytics of nonparental care (x)—the focus of our discussion—and thereby enables us to draw out lessons for the broader literature on childcare and child development.[36]

The choice of each form of care trades off the value of another hour of care for the child with the price of that care. The price of parental care is the forgone wage, w, whereas nonparental care has price per unit time, p. We assume p is small insofar as $w > p$ to account for the prevalence of informal, unpaid care, such as supervision by friends, grandparents, or older children (Yang et al., 2024). This calibration implies a relatively high opportunity cost of parental care time, m.

We may now consider how parental labor supply, n, responds to a shift in publicly provided supervision, g. The comparative static may be written as

$$\frac{\mathrm{d}n}{\mathrm{d}g} = \frac{z(\xi;\varphi)}{1 + z(\xi;\varphi)}, \tag{16}$$

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

$$z(\xi;\varphi) \equiv \frac{\left(\frac{\mu\xi}{1-\mu}\right)^{\varphi-1} + (1-\varphi)\xi^{-1} - \varphi}{\left(\frac{\mu\xi}{1-\mu}\right)^{1-\varphi} + (1-\varphi)\xi - \varphi}. \tag{17}$$

Derivations of all results in this section may be found in Online Appendix E.

The comparative static has two important properties. First, for $z > 0$, $\mathrm{d}n/\mathrm{d}g \in (0,1)$: labor supply falls when on-site time is reduced but less than one for one. While other labor supply outcomes are possible, the restriction $z > 0$ is a reasonable one. It obtains for any $\varphi \leq 0$ and, by continuity, over a range of φ to the right of zero. Indeed, equation (17) shows that if $\varphi > 0$, then $z > 0$ unless ξ is sufficiently small (the denominator is negative) or sufficiently large (the numerator is negative). Each of these polar cases conflicts with the ATUS evidence: a very small

[36] The Cobb-Douglas form has the awkward implication that $q \to 0$ as $g \to 0$. However, when paired with log utility, the scale of g has no allocative effect *via its role in* q. Rather, g shapes allocations through the time constraints.

(large) $\xi \equiv x/m$ implies that the marginal value of nonparental (parental) time is so low that m (x) responds far too much to a reduction in g (see Online Appendix E for a fuller discussion).[37]

Second, $\mathrm{d}n/\mathrm{d}g$ declines in φ (for any ξ). Therefore, at higher φ, labor supply falls *less* when on-site time is reduced.[38] Market work is sustained in this context by higher nonparental care. Intuitively, when the opportunity cost of parental time is high ($w > p$), a parent increases x relative to m if the two become more substitutable—that is, if the elasticity of substitution between them is increased. Hence, as φ is raised, a fall in g implies smaller increases in m, which require in turn smaller declines in n.

In light of our regression estimates, we assess the quantitative implications of equation (16) in the case where $\mathrm{d}n/\mathrm{d}g$ is small. We show that, in a neighborhood around $\mathrm{d}n/\mathrm{d}g = 0$, φ is bounded below such that $\varphi > (1+\xi)^{-1}$. To quantify the latter, we calibrate ξ to capture the initial allocation of childcare among parents "exposed" to school closure. For this purpose, we draw on Blau and Currie's (2006) figures for households where the mother had generally worked, which imply that children were under 1.36 hours of nonparental supervision per hour of parental care.[39] A value of $\xi = 1.36$ yields a *lower bound* of φ equal to 0.424. Thus, the pandemic-era data, as seen through this model, point to significant substitutability between forms of care.[40]

While this exercise aims to highlight the broader implications of our empirical results, one might be wary of generalizing from the pandemic period. For instance, 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

[37] We also show that a very low $\xi \equiv x/m$, and a very high m, emerge only under the alternative calibration $w < p$.
[38] When we vary φ, we adjust μ to hold fixed the initial value of ξ (and, thereby, n). See Online Appendix E for more.
[39] See primary and secondary arrangements in Blau and Currie's Table 2. The idea behind this approach is that households with employed mothers are arguably most "exposed" to a school closure. If the effect of a closure among them is nearly zero, then the average causal effect of policy will be nearly zero (as it appears to be, empirically).
[40] Note that $\mathrm{d}n/\mathrm{d}g \approx 0$ implies $\mathrm{d}m/\mathrm{d}g \approx 0$. The latter is consistent with estimates for the noncollege group.

more parental inputs (see Ramey and Ramey, 2010). Nevertheless, we see estimates in Berlinski et al. (2024) as broadly supportive of our conclusions. They study a sample of *pre*school children—a population for whom parental time is thought to be particularly crucial—and still find $\varphi = 0.92$ given a similar production function over parental and nonparental care.

The degree of substitutability between forms of care has significant implications for public policy and labor market dynamics. For instance, the price elasticity of demand for nonparental care increases in φ. Therefore, there will be greater take-up of subsidized care if nonparental time is highly substitutable for parental time.[41] Alternatively, consider a temporary increase in aggregate productivity that leads to higher wage offers. The Frisch elasticity of labor supply increases in φ: parents substitute more from childcare to market work if nonparental time is a close substitute for their own. See Online Appendix E for a fuller discussion.

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 a labor supply reaction. Even if we omit these controls, the labor supply responses represent a small fraction of the over 30 hours of childcare time "lost" with the suspension of in-person instruction. Time use data show that working from home while supervising children and nonparental private care helped support labor supply during school closures. The paper then integrates telework and nonparental care into a model of parental time allocation and illustrates how our results inform the identification of salient structural parameters.

[41] The federal government makes substantial investments in adolescent care. For instance, the Child Care and Development Fund made available $40 billion of subsidies to families of school-age children (U.S. Dept. of Health and Human Services, 2021).

Our empirical exploration of the roles of telework and nonparental care is limited, however, by the small sample sizes in the ATUS and by the paucity of direct measurements of nonparental care time.[42] We hope our work stimulates further efforts to measure these activities, which will in turn advance research into many related questions. To illustrate, these data would shed light on how shifts in the composition of the household—for instance, a grandparent or an older child moves in—alter the distribution of childcare and, therefore, parental labor supply.

With respect to the theory, we hope future research will push in two directions. One priority is to allow for more residential arrangements (e.g., two-parent households, the presence of a grandparent). This extension better captures the heterogeneity of care provider arrangements in the data (see Truskinovsky et al., 2022). A second priority is to model the link between on-site time and *specific* child outcomes, such as academic performance (see Jack et al., 2023, and Goldhaber et al., 2023, on test scores). This extension enables one to test if the predicted changes in parental time use patterns are consistent with evidence on academic outcomes.

[42] Surveys by Calarco et al. (2021) and Truskinovsky et al. (2022) are notable exceptions.

References

Aguiar, Mark, Erik Hurst, and Loukas Karabarbounis. 2012. "Recent Developments in the Economics of Time Use." *Annual Review of Economics,* 4:373–97.

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." *Economic Insights*, 8(4) 2–13.

Bacher-Hicks, Andrew, Tareena Musaddiq, Joshua Goodman, and Kevin Stange. 2024. "The Stickiness of Pandemic-Driven Disenrollment from Public Schools." *Economics of Education Review,* 100: 102527.

Barkowski, Scott, Joanne Song McLaughlin, and Yinlin Dai. 2024. "Young Children and Parents' Labor Supply During COVID-19." *American Journal of Health Economics*, forthcoming.

Barrero, José María, Nicholas Bloom, and Steven J. Davis. 2023. "The Evolution of Work from Home." *Journal of Economic Perspectives,* 37(4): 23–50.

Berlinski, Samuel, Maria Marta Ferreyra, Luca Flabbi, and Juan David Martin. 2024. "Child Care Markets, Parental Labor Supply and Child Development." *Journal of Political Economy*, 132(6): 2113–2177.

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(1): 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.

Cameron, Stephen V. and James J. Heckman. 2001. "The Dynamics of Educational Attainment for Black, Hispanic, and White Males." *Journal of Political Economy*, 109(3): 455–499.

Cascio, Elizabeth U. 2009. "Maternal Labor Supply and the Introduction of Kindergarten into American Public Schools." *Journal of Human Resources*, 44(1): 140–170. 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. 2024. "Time Use, College Attainment, and the Working-from-Home Revolution*." Journal of Population Economics*, 37(60).

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.

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 Public Economics*: 189: 104235.

Doepke, Matthias and Michèle Tertilt. 2016. "Families in Macroeconomics," in John B. Taylor and Harald Uhlig, eds., *Handbook of Macroeconomics*, Elsevier.

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.

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.

Flood, Sarah M., Liana C. Sayer, Daniel Backman, and Annie Chen. 2023. American Time Use Survey Data Extract Builder: Version 3.2 [dataset]. College Park, MD: University of Maryland and Minneapolis, MN: IPUMS, 2023. https://doi.org/10.18128/D060.V3.2

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. 2024. "Childcare Responsibilities and Parental Labor Market Outcomes During the COVID-19 Pandemic," *Journal of Labor Research*, 45: 153–200.

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.

Goldhaber, Dan, Thomas J. Kane, Andrew McEachin, Emily Morton, Tyler Patterson, and Douglas O. Staiger. 2023. "The Educational Consequences of Remote and Hybrid Instruction During the Pandemic." *American Economic Review: Insights* 5(3): 377–392**.**

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.

Hansen, Benjamin, Joseph Sabia, and Jessamyn Schaller. 2024. "Schools, Job Flexibility, and Married Women's Labor Supply." *Journal of Human Resources,* forthcoming.

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 COVID-19 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.

Jack, Rebecca, Clare Halloran, James Okun, and Emily Oster. 2023. "Pandemic Schooling Mode and Student Test Scores: Evidence from US School Districts." *American Economic Review: Insights* 5(2): 173–90**.**

Kane, Thomas J. 1994. "College Entry by Blacks Since 1970: The Role of College Costs, Family Background, and the Returns to Education." *Journal of Political Economy*, 102(5): 878–911.

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

Kwon, Sarah Jiyoon. 2024. "Grandparents and Parental Labor Supply During the COVID-19 Pandemic." *Review of Economics of the Household*, 22: 935–964.

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.

Mulligan, Casey B. 2001. "Aggregate Implications of Indivisible Labor." *Advances in Macroeconomics*, 1(1).

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(2): 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.

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

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

Truskinovsky, Yulya, Jessica M. Finlay, and Lindsay C. Kobayashi. 2022. "Caregiving in a Pandemic: COVID-19 and the Well-Being of Family Caregivers 55+ in the United States." *Medical Care Research and Review*, 79(5): 663–675.

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. 2025. "Patterns in Receiving Informal Help with Childcare Among U.S. Parents During the COVID-19 Pandemic." *Journal of Family Issues*, 46(1): 230–269.