New Goods and Labor Hours: Evidence From A Natural Experiment.

Paul Scanlon *

Trinity College, Dublin.

Abstract

This paper incorporates new goods into a standard model of labor-leisure choice, and explores how greater product variety affects labor supply. According to the model, the introduction of a new good increases marginal utility and in turn the incentive to earn income. To evaluate this mechanism, I exploit the staggered introduction of a significant new product—the Apple iPhone—to four U.S. states. For young part-time workers, I find that labor hours rose by approximately two hours per week in the weeks surrounding an iPhone introduction.

JEL Codes: E20, E21, J20, J22.

Keywords: Labor Supply, Innovation, Technology.

*I thank numerous seminar participants for helpful comments and Niall Maher who provided excellent research assistance. Email: scanlop@tcd.ie

1 Introduction

Previous empirical work in labor economics has emphasized the role of the wage in determining labor hours. In this paper, I focus on the role of consumption and specifically on how increasing product variety affects labor hours. I present a model where labor supply is a function of product variety and examine how the introduction of new goods affects marginal utility. The concept of a love of variety has been fruitful in explaining an array of economic phenomena, and in the context of the labor market, provides a natural motivation to earn income. Because product development is a central concern of firms and underlies rising living standards, determining how consumers respond has important economic implications. Particularly for analysing the incentives to work, failing to account for this channel can lead to an underestimation of predicted labor hours.

According to the standard labor-leisure model, people supply labor to finance the consumption of a single good. With a dominant income effect, a permanent rise in the wage translates into more consumption and this causes marginal utility to decline, reducing labor hours. Yet in reality the nature of consumption changes over time, affecting marginal utility and the desire to earn income. By increasing marginal utility, more product variety attenuates the income effect of a rising real wage and raises the incentive to supply labor. Analogous to the way technology offsets diminishing returns to capital, expanding product variety counters diminishing marginal utility and motivates higher labor hours.

To evaluate the link between labor hours and new goods, I examine how the staggered introduction of a significant new product—the Apple iPhone—affected labor hours of young, non-prime workers in the United States. Starting from its introduction in 2007, AT&T had exclusive rights to sell the iPhone. Yet absent AT&T coverage, the iPhone was initially unavailable in four U.S. states: Montana, North and

South Dakota, and Wyoming. In early 2011, however, Apple and another network, Verizon, introduced the Verizon iPhone, which became available nationwide. Because the goal of both firms was to increase market share across all states, the timing of this alliance was independent of local economic conditions. Exploiting this event and treating it as a plausibly exogenous increase in variety, I examine how labor hours changed in the foregoing states around this time. Facilitating identification is the differential impact the product introduction had across states in early 2011, coupled with its varying appeal across age groups. For young non-prime part-time workers in the four affected states, I find that the introduction of the Verizon iPhone accounts for an approximate 10 percent rise in hours per week in the weeks surrounding the release date.

I proceed as follows. Section 1 outlines a model showing how increasing product variety affects labor hours. In Section 2, I investigate how the introduction of the Verizon iPhone affected labor hours in four U.S. states and find that the model's main prediction is borne out by the data. To my knowledge, it is the first paper to document how the introduction of a significant new product affects labor supply. Finally, Section 3 concludes.

2 Model of Labor Hours and Variety

The empirical strategy of Section 3 informs the theoretical model presented here. To sharpen identification, the empirical section focusses on young part-time workers and explores how a new durable good affects labor hours. Compared to full-time prime-aged workers, these are more likely to be liquidity constrained, implying that the effects on labor supply are concentrated over a short period. Motivated by this, I present a stylized model examining a new product introduction in the presence of a liquidity constraint. Yet this constraint only impacts the distribution of labor hours and has no bearing on the quantitative importance of the underlying mechanism.

2.0.1 The Economic Environment

The consumer lives for two periods. Utility is separable and derives from nondurable consumption c, the stock of durable goods D, and labor supply l. There is no discounting, the gross interest rate is 1, and the wage w is constant across periods. The time endowment each period is one. Lifetime utility is

$$U = u(c_t) + u(D_t) - v(l_t) + u(c_{t+1}) + u(D_{t+1}) - v(l_{t+1}),$$

where $u' > 0$, $u'' < 0$, $v' > 0$, $v'' > 0$, and $u(D) = 0$ if $D = 0$. A new durable good is introduced in period t, and the evolution of the durable good stock is

$$D_{t+1} = (1 - \delta)D_t + e_t,$$

where δ denotes the durable stock depreciation rate and e_t the expenditure flow on durable goods.[1] Noting that the relative price of the durable good is p_t, the intertemporal budget constraint is

$$c_t + c_{t+1} + p_t e_t = wl_t + wl_{t+1}. \quad (1)$$

There is a borrowing constraint in period t implying that

$$c_t + p_t e_t \leq wl_t. \quad (2)$$

2.0.2 Solution

It is convenient to set $p_t = 1$ and given that the durable good is new, $D_t = 0$. It follows that the Kuhn-Tucker optimality conditions for labor hours are

[1] For convenience, I ignore indivisibilities and the non-integral nature of the variable D.

$$w\lambda = v'(l_{t+1}),$$

where λ is the Lagrangian multiplier on Eq. 1, and

$$w(\lambda + \mu) = v'(l_t),$$

where μ is the Lagrangian multiplier on (2). Combining, the optimal allocation of labor hours across both periods is implicitly given by

$$\frac{v'(l_t)}{v'(l_{t+1})} = 1 + \frac{\mu}{\lambda}.$$

With a binding borrowing constraint, $\mu > 0$, implying that $l_t > l_{t+1}$.[2] Thus the model predicts that the introduction of a new durable good induces higher labor hours in period t. Because of the new durable good, there is a greater marginal utility gain to working in period t, and absent the ability to borrow, this motivates higher labor hours. In the next section, I examine this prediction empirically. Without a borrowing constraint, labor hours would be higher and equal each period.[3] Absent the new durable good, labor hours would be equal each period, but at a lower level.

3 Empirical Framework: Testing the Model

According to the model, greater product variety causes labor hours to rise. Yet a number of factors complicate identification in aggregate data. As an example, increases in product variety are often coincident with more product market deregulation and

[2]To see why the constraint binds, suppose that it is nonbinding and thus $\mu = 0$. Then the optimality conditions imply that $l_t = l_{t+1}$ and $c_t = c_{t+1} = e_t$. Yet these equalities contradict the constraints (1) and (2). This implies that $\mu > 0$. Intuitively, given the two products to purchase in period t, there is a greater marginal utility gain to working. Also, given nonseparability and the borrowing constraint, in period $t + 1$ the levels of consumption and hence labor hours are unaffected by the new good.

[3]Empirically, this makes the effects difficult to identify, particularly for repayments scheduled over many years or decades.

hence greater labor demand. Meanwhile, at longer horizons increases in labor supply raises market size and induces greater innovation. In both cases, increases in labor supply accompany greater variety growth, but through channels other than the one emphasized here.

Given these concerns, I turn to less aggregated data. What is required to test the main mechanism is i) the introduction of a specific new product; and ii) workers with limited access to funding, but with flexible labor hours, who seek to purchase that product. While a weakness of this approach is it only examines a single good, it nonetheless provides evidence on whether the mechanism is potentially of quantitative importance.

One clear-cut example of such a good is the Apple iPhone. Being a popular innovation with a high relative price, the theory predicts that its release would lead to a rise in labor supply for young liquidity-constrained workers. For identification, I exploit both time series and cross-sectional variation in iPhone introduction across U.S. states. This variation arises primarily from the late introduction of the iPhone to four U.S. states in February 2011: Montana, North Dakota, South Dakota, and Wyoming. I classify these as the treatment states. Additional variation comes from the disparate appeal of the iPhone across different age cohorts.

3.1 Event Setting and Treatment Group

From its initial release in June 2007 until January 2011, the U.S. wireless carrier, AT&T, had exclusive rights to sell the iPhone, bundled with their wireless contracts. Yet absent AT&T coverage over this period, the treatment states were without iPhone connectivity. Starting in January 2011, however, Verizon, a rival of A&T and the largest U.S. phone carrier with presence in all states, secured rights to sell iPhone activations. The official announcement of the Verizon iPhone occurred on 11 January 2011, and it became available for sale on 10 February. Especially relevant here, this introduction was related to a nationwide contract between Verizon and Apple and was incidental to developments in these states. Depending on the precise model, there was an upfront cost of either $199 and $299 to purchase the phone, a significant amount relative to a $9 hourly wage. Given that this price was subsidized, it entailed a commitment to a two-year contract and data plan, entailing payments of approximately $30 dollars per month. Examining how labor supply changed across states around the introduction time therefore provides a useful opportunity to test the theory. Significantly, the period of analysis, 2010–2011, is a period characterized by high public consciousness of the iPhone. Contrasting with the 2007–2009 period, where there were a combined 16 million activations, there were 15 million U.S. activations in 2010 and 30 million in 2011. Between the product release in February 2011 and the end of March 2011, Verizon activated 2.3 million iPhones.

Throughout, I focus on workers with flexible labor hours for whom the iPhone offers particular appeal. One consumer segment satisfying these criteria are non-prime part-time workers in the 15–24 age category, whom I designate the "treatment group." Offering access to gaming, music, shopping, and a variety of social networks, the iPhone appeals on multiple dimensions to this market segment. For example, a survey conducted by a financial services company, *Piper Jaffray*, in Spring 2010 revealed that 31 percent of teenagers planned to purchase the iPhone over the following six months—a rise from 16 percent a year earlier.[4] Because many in this age group sought a new iPhone, but did not have one, the iPhone represented a significant new product at this stage. According to the website *Statista*, in July 2011 the 13–24 age group owned 23 percent of U.S. iPhones; by comparison, the figure for those aged 55 and over was only 14 percent. Focusing on a specific age group with

[4]According to the bi-annual survey, "*Taking Stock with Teens*," 15 percent of teenagers owned an iPhone in Spring 2010, compared to 48 percent three years later. This suggests that product diffusion was relatively low among younger people up until Spring 2010, but increased markedly thereafter. The survey sample comprises approximately 8,000 high-school students from middle to high income households. By Spring 2019, 83 percent of teenagers owned iPhones.

similar wages ensures that preferences are relatively homogeneous across workers, resulting in a sharper test of the mechanism.

In addition to their interest in purchasing the product, the financial circumstances of the treatment group also sharpens identification. This age cohort has fewer financial commitments, making it easier to isolate movements in labor supply attributable to exogenous changes in the desire to work. Having few assets and lacking stable income and collateral, they are also likely to have limited access to loans. Faced with such liquidity constraints, labor supply provides an important source of income. Among 15–24 year-olds in the treatment states, around 50 percent engaged in some form of part-time employment throughout 2010 and 2011, where they faced an average weekly wage of approximately nine dollars an hour. Costing between \$199 and \$299, the iPhone was therefore a significant purchase, requiring around 28 hours of labor. An additional advantage of focusing on part-time workers is their versatility in adapting to a variety of work schedules.

Another factor facilitating identification is the occurrence of a significant share of sales during the quarter of introduction. One reason for this is the publicity generated by a new iPhone announcement and release. Reports of large lines and depleted inventories routinely accompany new releases—motivating purchases in the period around the release date. Amplifying this trend are demonstration, peer, and network effects, occurring soon after the release date. Compared to older cohorts, younger cohorts are also more likely to have shorter planning horizons, raising the likelihood of a labor response around the release date.

To systematically gauge the degree of interest in the iPhone across states and over time, I use Google Trends data. This provides information on the proportion of statewide Google searches seeking information on a specific search topic. Most importantly, Choi and Varian (2012) report that such product searches predict future sales activity. Using this data, Figure 1 displays U.S. search interest for "iPhone"

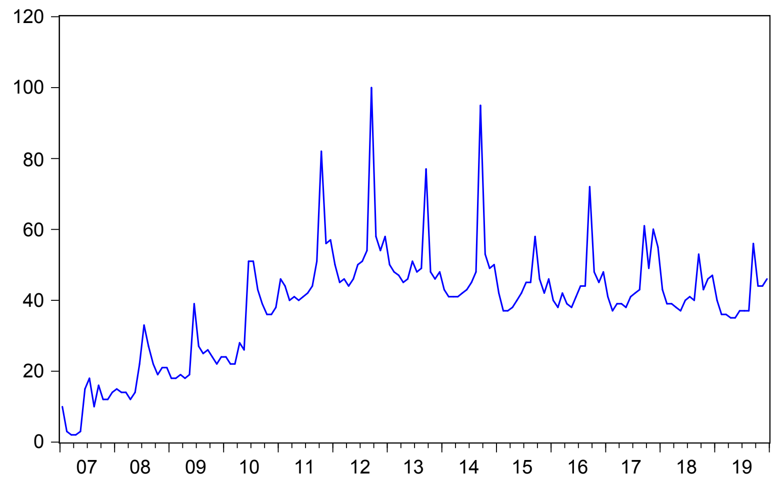

Figure 1: GOOGLE SEARCHES FOR "IPHONE" : U.S., 2007-2019
Source: Google Trends

between 2007-2019. More specific to our concern, Figure 2 contrasts the degrees of interest over time across the treatment and control states.[5] Figure 2a depicts an increasing search intensity across the four treatment states that was almost double the level shown for the June 2010 release. In particular, January and February 2011—the announcement and release months—marked the highest degree of search intensity across the treatment states. By contrast, these months were accompanied by a relatively small deviation from the upward trend in the iPhone search series for the control states. In addition to basic Google searches, I also track prospective buyer interest through YouTube searches. Because such searches are more likely to arise from younger cohorts, they convey a more accurate picture of the interest shown by younger cohorts. Apparent from Figure 3 is the significant search intensity across the treatment states between from the announcement date on 11 January until the end of February. Together these observations underscore the greater significance of the Verizon iPhone across the treatment states.

Another consideration is to what extent the early 2011 introduction was expected.

[5]As discussed in Section 3.3, the controls comprise all states bordering the treatment states and the remaining Midwestern states. The indices weigh search interest by statewide population.

For many years up until mid–2010, there were repeated rumors that Verizon would attain network rights to sell the iPhone, none of which materialized. Yet, in what marked a departure from previous media reports, an article in the *Wall Street Journal* on October 7, 2010 revealed detailed information about actual production plans and an early 2011 release date. The *New York Times* published a similar article on October 10.[6] Following these reports, search interest in the "Verizon iPhone" intensified. Evident from Figure 4 is a rising search interest throughout the US from mid-October, which was unusually sustained, and culminated with a pronounced rise in December. Focusing on the cross section, Figure 5a shows a relatively high search interest across the treatment states over the period October to December 2010. Here, darker shades signify a greater search intensity, *relative* to all Google searches in the specific states. Testifying to latent interest in the product, Figure 5b shows how the release of the iPhone 4 in June 2010 also stimulated interest in a prospective Verizon iPhone across the treatment states.

3.2 Data

Data on labor hours are from the Current Population Survey (CPS) Basic Monthly files, downloaded from the Integrated Public Use Microdata Series (IPUMS). To maintain data comparability, I start the analysis following a CPS revision in 1994. The data spans the 26-year period, 1994–2019. The survey covers a representative sample of around 60,000 households each month from across all U.S. states. Each household participates in the survey for four months, then departs for eight, and returns for another four the following year. As a result, there is a staggered introduction of households into the survey: each month a quarter of the households leave and are replaced by an incoming cohort. To calculate average wages, I use data from the CPS

[6] The respective articles were "*Apple Readies Verizon iPhone*" by Yukari Iwatani Kane and Ting-I Tsai on 7 October 2010 and "*Apple to Offer iPhone on Verizon, Ending Exclusivity with AT&T*" by Miguel Helft on 10 October, 2010.

Merged Outgoing Rotation Groups (MORG) files available from IPUMS. In contrast to the Basic Monthly CPS files, this survey asks departing households—i.e., those in months four or eight—to specify their wages in the previous week.

Because of the shortness of the panel dimension and the high rates of attrition for younger people, I treat the data as repeated cross sections. The main variable of interest is reported hours worked last week at all jobs for non-prime part-time workers aged between 15-24. On the week including the 19th day of each month, the head householder reports the labor hours of each household member during the week of the 12th of that month. To exclude outliers, I restrict the analysis to those working more than five hours per week, and drop those working more than 60 hours per week.[7] As a measure of household income, I use the family income measure from the CPS, deflated by the Consumer Price Index. This categorizes all households into income bins according to total family income from all sources in the previous year. Because there is an open income range for incomes exceeding $\$150,000$, I exclude households from this highest income bracket—a segment comprising four percent of the treatment group and one likely to have little treatment exposure. Over the treatment period, there are 627 observations remaining for the treatment group. Although modest in size, the treatment sample covers a specific, relatively homogeneous segment of the labor market engaged in work with similar compensation and job requirements.

3.3 Differences-in-Differences Setup: Experimental Design

To identify the effect of the Verizon iPhone introduction on labor hours, I use a differences-in-differences framework. With the Dakotas, Montana, and Wyoming comprising the treatment states, the main event I consider is the introduction of the

[7] Included in the sample are those who usually work part-time, but work more than 35 hours in a particular sample week.

Verizon iPhone in early 2011. Because the treatment states exhibit considerable homogeneity and are less susceptible to fluctuations in economic activity—such as tourism—than other states, they are especially conducive to identifying the effects of an exogenous shock. Identification comes from comparing labor hours across a number of dimensions: i) before, during, and after the treatment period for younger workers in the treatment states; ii) between younger workers in the treatment and control states; and iii) between younger and older workers in the treatment states. The treatment months comprise the announcement and release months, January and February 2011, together with December 2010 to capture an anticipation effect. Because most of the treatment occurs during the off-season period of January and February, the likelihood of confounding factors affecting labor hours is relatively low.

Given sampling variability associated with younger workers, I choose as large a set of control states as possible. With the focus on part-time work, it is particularly important to choose states exhibiting similar seasonal patterns and hence of similar latitude. Figure 7, from Peterson (2011), illustrates the extent of wireless AT&T coverage across U.S. states in January 2011. Given that the iPhone 4 operated on 3G and Edge cellular networks, the lack of coverage across the treatment states is striking. Also notable is the relatively poor AT&T coverage across rural areas of many states. To ensure minimum treatment exposure, I therefore use as controls the metropolitan areas of i) states bordering the treatment states; and ii) all remaining Midwestern States.[8] Because the timing of the iPhone introduction was not endogenous to time-varying determinants of labor hours in the treatment states during the treatment period, the main identifying assumption is parallel trends in pretreatment labor hours. The fact that young non-prime part-time workers engage in similar work and face similar wages suggests that this assumption is reasonable. Table 1 displays summary statistics for the treatment and control sample, indicating

[8]The control states consist of Idaho, Utah, Colorado, Nebraska, Minnesota, Iowa, Wisconsin, Illinois, Indiana, Ohio, Michigan, Missouri, and Kansas.

that both groups are broadly comparable across a number of relevant characteristics.

One factor that could lead to an understatement of the treatment effect is people in the control regions purchasing the Verizon iPhone to replace their existing model. Yet three factors mitigate this concern. First, an almost identical iPhone 4 had been released by AT&T in June 2010, substantially reducing the appeal of the Verizon iPhone. Second and related, by January 2011 there was an incentive to wait for the next iPhone release, which typically took place during the summer. Third, those on existing AT&T contracts faced high switching costs of 325 dollars if switching provider. For these reasons, switching offered little appeal to less affluent younger workers in the control states.

3.4 Regression Specification and Results

To capture seasonality and the fact that average labor hours are lower in some states, I include month, year, and state fixed effects, together with state-month and state-year interactions in the regression. The interaction terms capture factors at the state level, which could affect statewide demand or supply for labor in specific months or years. To correct for serial correlation across individual states over time, I cluster standard errors at the state level. The main regression specification is

$$h_{ist} = \alpha + \delta_i + \tau_{yt} + \xi_{mt} + \zeta T_t + \tau_{yt} \times \delta_i + \xi_{mt} \times \delta_i + \beta(T_t \times D_i) + X_{ist}^{'}\daleth + \epsilon_{ist}, \quad (3)$$

where h_{ist} denotes weekly hours by part-time worker i in state s at time t, α is an intercept, δ_i a state fixed effect, τ_{yt} a year fixed effect, ξ_{mt} a month fixed effect, X_{it} a vector of covariates, and ϵ_{ist} an error term. T is an index, indicating the presence of treatment and D is a dummy for the treatment states. Because it is an important variable in determining consumption and labor hours of younger household members, I include household income as a control. To account for time-varying economic conditions across states and to increase the precision of the estimation, I also include the monthly state unemployment rate. Because there is no reason to expect a uniform treatment effect across personal attributes—such as age—controlling for such characteristics makes less appeal in this context. For example, the product could attract people of a certain age into the workforce, thereby changing the age distribution of workers and making age an outcome variable. For robustness, I also present the analysis with and without such controls.[9] For ease of exposition, I use actual as opposed to log hours.

Examining Google Trends data in Figure 2a, the search intensity is highest in January and February. To capture this, I set a lower treatment intensity of .5 in December to capture an anticipation effect and 1 in the announcement and release months. The main coefficient of interest is β, the interaction of the treatment states and months. Positive and significant values indicate higher relative weekly labor hours for the average young non-prime part-time worker in the Dakotas, Montana, and Wyoming during the treatment months.

3.4.1 Assessing Parallel Trends

To evaluate the assumption of parallel trends, I take an event study approach and apply the treatment each month for eight months before and after the treatment period; i.e., from April 2010 to October 2011:

$$h_{ist} = \alpha + \delta_i + \tau_{yt} + \xi_{mt} + \zeta T_t + \xi_m \times \delta_i + \gamma_i t + \sum_{t=-9}^{t=11} \beta_t (T_t \times D_i) + X'_{ist} \daleth + \epsilon_{ist}, \quad (4)$$

where t denotes a linear trend term, γ_i the state-specific trend coefficient, and β_0, β_1,

[9]Data limitations prohibit the inclusion of wages in the main regressions. Wages were relatively stable over the short treatment period.

and β_2 the treatment effects for December 2010–February 2011. I normalize the treatment effect to zero in November 2010, and this constitutes the baseline month. I pool all months prior to April 2010 and all months after October 2011 and include them as separate dummies, T_{-9} and T_{11}. Because the objective is to assess the assumption of pretreatment parallel trends, to preserve efficiency I replace state-year fixed effects with state-specific linear trends.

Figure 8a plots the treatment coefficients, along with the associated 95 percent confidence intervals. Three points stand out. First, in the pre-treatment period, the confidence intervals for the treatment effects center around zero, and the effects are individually and jointly insignificant at the 5 percent level. This absence of pretrends between treatment and control states lends support to the validity of the control group.

Second and most important is the marked divergence from the preexisting pattern, commencing in December 2010. The treatment months are the only three consecutive months whose confidence intervals lie above zero and where each exhibits a statistically significant treatment effect. Especially given it is likely harder to raise hours in the predominantly less urbanised treatment states, the sharp and discontinuous rise in hours strongly suggests the presence of a treatment effect. Coincident with the largest spike in the Google Trends series, the most pronounced effect is for the announcement month, January, where mean hours rise most. Also notable is the effect diminishes gradually. This persistence could be reconciled with the theory if workers increased hours over a number of months or if the product became more appealing as network effects developed. To increase power, Figure 8b excludes the bottom 30 percent of workers by family income and thus focuses on those for whom this relatively expensive product may have more appeal. Evident from the figure is a stronger and more precisely estimated treatment effect for this segment, particularly for January 2011.

Estimating distinct treatment effects across many months places considerable demands on the data, leading to wide confidence intervals in Figure 8. Because the focus of the analysis is the total treatment effect, to enhance power I break the year into three-month intervals. These intervals correspond to natural seasonal categories and facilitate a more precise analysis of trends over a longer period. The quarters comprise i) December, January, and February; ii) March, April, and May; iii) June, July, and August; and iv) September, October, and November. I then perform the event study analysis over three years for the period commencing June 2009 and ending August 2012, and create separate dummies for the periods before and after these months. The baseline is the three-month period starting September 2010. Figure 9 displays these more precisely estimated treatment effects. Similar to Figure 8, it shows a pronounced rise in hours during the treatment period, with a stronger effect when excluding the bottom 30 percent of workers by family income.

3.4.2 Main Experiment: Verizon iPhone Introduction

Table 2 summarizes results for the 2011 experiment for the baseline 15–24 age group and other demographic subgroups. For convenience, I omit all fixed effects and controls from the regression tables. Column 1 presents the baseline regression results for hours worked per part-time worker under the age of 25. This coefficient addresses the question: conditional on working, what are the expected additional labor hours per worker per week in the treatment relative to the control states during the treatment months? Conditional on all of the controls, this value is 2.3 hours when the treatment intensity is one (i.e., in January and February), and this corresponds to around a 10 percent rise relative to average hours. The associated *t*-value is 5.8, and the 99 percent confidence interval comprises only positive values, providing strong evidence for an economically significant increase in labor hours. Aggregated over the treatment period, the rise in labor hours leads to additional income of approximately

200 dollars—approximately the cost of the iPhone. Across all age groups, the greatest response is from the 18–21 cohort.

To sharpen identification, Columns 4-7 display results for broader demographic subsamples. Across all age groups, there is a positive, but insignificant treatment effect. Column 4 restricts the regression to part-time workers aged 25 and over. Because this cohort has greater access to funds and is more likely to have an existing Android phone, they face less treatment exposure. As shown, the treatment effect for this group is insignificant. In particular, this indicates that the results in Column 1 are unrelated to a general shift in the demand for part-time workers.

Yet the aggregate effect for those aged 25 and over masks considerable heterogeneity. While there is no treatment effect for those aged 40 and over, Column 5 reveals a significant and negative effect for those aged between 25 and 39. To illustrate the divergent trends, Figure 10a plots the treatment coefficients from Eq. 4, where I restrict the difference-in-differences analysis to the 25–39 age group across treatment and control states. This highlights the negative treatment effect. To shed further light on this, Figure 10b plots the treatment coefficients, where the treatment group is the 25-39 cohort, and the control group is the 15-24 cohort, each from the treatment states[10] Presented in this way, it is clear that both cohorts within the treatment states exhibit parallel pre-trends, but then diverge sharply over the treatment period. One natural interpretation of these divergent movements is they reflect a general equilibrium effect whereby the younger cohort sought more labor hours over the treatment period and displaced similar workers in the 25–39 category.[11] Compared to older cohorts, the 25–39 segment is more substitutable for younger workers and is more likely to work in similar settings. Absent wage adjustment, this reallocation is

[10]To isolate the effects of the seasonal summer market, I include interaction effects for the three-month period, June, July, and August with each of the years 2010 and 2011.

[11]Smith (2012) documents increased competition among adults and teenagers for low-skill service jobs.

an implication of any model where output—and hence labor—is demand-determined in the short run.

To show formally that the product had a differential impact by age, Figure 11 plots treatment coefficients from a triple-difference specification, using part-time workers aged between 25-39 in the treatment states as an additional control. As already noted, these workers had fewer financial constraints and were more likely to have an existing phone, and thus faced lower treatment exposure. For this specification, there is a treatment effect of 3 with t-value of 6.7. Given the evidence above that hours fell for the 25-39 cohort in the treatment states, the triple interaction effect is unsurprisingly highly significant and large. Yet because the labor hours of one control group—the old in the treatment states—plausibly depend on the hours of the treatment group over short horizons, this result is only suggestive.

Table 3 shows the results by education level. Possibly reflecting their greater flexibility and work experience, the strongest effect is for those outside of education, a group comprising 35 percent of the treatment group. Also displayed are the results excluding college students, a group comprising 37 percent of the treatment group. One issue with college students, however, is they appear as resident in the home state, but may be attending college outside of that state. Although this could contaminate results, the treatment states have a net inflow of students, suggesting that this issue is unlikely to be quantitatively important. Nonetheless, there is a stronger treatment effect when excluding college students. Another explanation is students had less hours flexibility.

Next I investigate the role of family income, shown in Table 4. Insofar as social preferences vary across income levels, the iPhone may have more appeal as income rises. In addition, higher family income provides potential access to family plans, reducing the effective cost of an iPhone. While these channels raise hours, the greater likelihood that high-income parents gift other family members an iPhone reduces hours. Taken together, this suggests that family income enters nonlinearly, and the data bear this out: at low incomes, there is little treatment effect, but as income rises, labor hours increase, but then decline again in the top decile. As already shown in Figures 8 and 9, there is a stronger treatment effect when excluding workers from the bottom 30 percent of the household income distribution.

I devote the remainder of this section to robustness checks. For comparison, Column 1 of Table 6 displays the baseline results. One concern is the rise in hours reflects a composition effect of many people with a high taste for work entering the survey and remaining there for multiple periods. To mitigate this issue, I restrict the regression to those who appear only once over the treatment period.[12] While this does not preclude a composition effect, results strengthen for this subsample, suggesting that this particular channel is not driving results. Columns 3–6 display the results when excluding each of the treatment states. Results are generally similar, indicating that the treatment effects occur broadly across the treatment group. Finally, as a falsification test (not shown), I apply a placebo treatment to December, January, and February of each year from 1994 onwards, and find that the only three-month period exhibiting a significant effect at the 5 percent level is the actual treatment period. The next highest t-statistic is 1.70 for the three-month period commencing December 2012, where the attendant treatment coefficient is 0.54. All remaining years are insignificant.

Table 5 includes gender, age, eduction, wage, and industry as additional controls. Including age reduces the treatment effect, indicating an age-dependent treatment effect. Consistent with this, and as revealed in Table 2, there is a pronounced rise in hours among the 18-21 age group. Throughout, the state unemployment rate remains insignificant, suggesting that local economic conditions have a limited bearing on this segment of the labor market. Because data on wages are only available for respondents in their fourth or eight survey months, their inclusion leads to a 70

[12]This subsample comprises survey respondents in weeks 1, 4, 5, and 8 of their eight-month survey schedule.

percent reduction in the treatment sample. For these respondents, who appear once in the treatment period, higher wages are associated with greater labor hours, while the treatment effect rises.

To determine the sensitivity of results to control selection, Table 7 displays results using five alternative sets of controls: i) all areas of the control states; ii) suburban areas of the control states; iii) rural areas of the control states; iv) all areas of the six bordering states; and v) rural areas of states where AT&T coverage was extensive (Illinois, Indiana, and Michigan). As revealed in Figure 7, many rural areas were poorly served by AT&T and thus may be susceptible to treatment. Reflecting this possible exposure to treatment, results are weaker and less precisely estimated when rural areas comprise part of the control group. Comparing columns, there is also a modest reduction in the treatment effect for the less urbanized bordering states, yet it remains highly significant with a t-statistic of 4. More interesting is the final column, where the controls comprise the rural areas of states that had widespread coverage. These rural regions likely have similar underlying fundamentals to the more rural treatment states, and therefore represent a particularly useful benchmark. The treatment effect is significant and of comparable magnitude to the baseline, but is less precisely estimated. In summary, regardless of the control group, results remain significant throughout.

3.4.3 Natural Experiment 2: iPhone 4 Release in June 2010

Next I explore how the introduction of the iPhone 4 affected labor hours across the treatment and control states.[13] Announced on 7 June, 2010 and released on 21 June, iPhones were only operational across the control states at this stage. In this "reverse experiment," the model predicts a negative treatment effect in the Dakotas, Montana,

[13]Survey evidence from *Piper Jaffray* indicates little penetration of the iPhone among teenagers up until 2009, with only 7 percent owning an iPhone by Spring 2009.

and Wyoming; i.e., higher relative hours in the control states. Figures 2a and 2b show the contrasting search intensity for the “iPhone” around the release date. Focusing on searches for the “iPhone 4”, Figure 11b highlights the relatively small degree of interest across the Dakotas, Montana, and Wyoming in June and July 2010.

In contrast to previous summer releases, a number of factors make 2010 a particularly useful setting to test the model. Owing to its introduction after the Great Recession, there was a lower likelihood of confounding business cycle factors at this stage. As highlighted in Figure 1 and discussed in Section 3.1, there was also significant interest in iPhones by this stage. Finally, unlike releases in 2008 and 2009, which were extensions of the original iPhone, the iPhone 4 represented a substantially new model. Nonetheless, the summer labor market is associated with different labor market dynamics, making the summer release a less useful test than the first event. First, a number of factors—such as work experience or internships—motivate entry into the labor force over the summer, leading to greater labor hours volatility. Contributing to this volatility is the vibrant and flexible summer services industry, affording the opportunity to work more shifts of varying duration. Given this greater likelihood of confounding factors and compositional changes, my preferred experiment remains the former one.

For the analysis, I set the treatment intensity to 1 in June and July 2010 and 0 otherwise. Including July ensures that the treatment period includes a survey month after the release date. Table 8 displays the results for this event. The treatment effect is -1.23 and highly statistically significant, yet smaller in magnitude and less precisely estimated that before. Column 3 displays results from a regression incorporating both experiments, showing that each estimate is of the predicted sign. While not pronounced, Figure 8a shows a mild deviation from an approximate zero trend line in June and July 2010.

3.5 Review

The empirical findings are subject to a number of caveats. For the first experiment, the announcement and launch weeks coincide with the survey weeks, and this could lead to an overestimation of the weekly treatment effect. Countering this is the anticipated arrival of another release during the summer months, reducing the incentive to purchase the Verizon iPhone in early 2011.

While the analysis shows that new goods *can* affect labor hours, a number of factors potentially compromise external validity in this setting. The first is the nature of the product, which is a significant innovation, and one subject to peer and network effects. A related issue is the study comprises only young part-time workers—a cohort frequently associated with such peer effects, and one unrepresentative of the general population. Yet two factors are worth noting here. First, Luttmer (2005) documents the presence of interpersonal preferences across all age groups, suggesting that peer effects occur more broadly.[14] Moreover, while features specific to the natural experiment—e.g., the product itself and the age cohort—facilitate identification and make the main channel transparent, they plausibly have greater impact on the timing and magnitude of the change in hours, rather than the change itself. In particular, they do not undercut the main channel whereby a new product raises marginal utility and in turn the incentive raise labor hours. It would be surprising, for instance, if this impulse was operative for a 22-year old, but absent upon becoming a full-time employee.[15]

Despite the above concerns, the results nonetheless confer an interesting insight into what motivates young people to enter the labor market. Ashworth et al. (2017)

[14] This channel would amplify the role of new goods as a motivation for labor supply. In his book, *Luxury Fever*, Frank (1999) argues that interpersonal comparisons are important drivers of luxury consumption expenditure and household debt. In her book, *The Overworked American*, Schor (1993) outlines a similar narrative and highlights the interaction between labor hours and consumerism.

[15] Although speculative, the data are consistent with an equilibrium whereby the continual arrival of new goods attenuates the incentive to reduce labor hours below the standard 40–hour work week.

document the long-run benefits to work experience and find that the return to an additional year of experience exceeds the return to an additional year of school or college. Among other benefits, work experience provides on-the-job training, promotes greater greater labor market attachment in later years, and fosters qualities such as punctuality and discipline. Determining what motivates younger people to work is therefore an important topic in its own right.

Another setting where the finding may be especially relevant is developing countries. In these economies, self-employment, part-time work, and financial constraints are commonplace, and the findings here suggest that opening up markets—e.g., through trade liberalization—can potentially lead to greater labor market participation. Anecdotal evidence documented by Berg (1961) provides some support for this channel in developing countries.

4 Conclusion

This paper presents a model that highlights how expanding variety growth can raise labor hours. By raising marginal utility, greater product variety increases consumption demand and, in turn, the incentive to supply labor. An attractive feature of the framework is it departs only minimally from the standard model, yet it can reconcile a dominant income effect and wage growth with non-declining labor hours. Examining the staggered introduction of the Verizon iPhone to four U.S. states, I evaluate the model's prediction that a rise in product variety raises labor hours. Consistent with the main mechanism, I document a significant rise in hours by young part-time workers in the weeks surrounding the introduction date. An interesting path for future work would be to examine how general this finding is.

References

Ashworth, J., J. Hotz, A. Maurel, and T. Ransom (2017). Changes across Cohorts in Wage Returns to Schooling and Early Work Experiences. NBER Working Papers 24160, National Bureau of Economic Research, Inc.

Berg, E. J. (1961). Backward-Sloping Labor Supply Functions in Dual Economies - The Africa Case. *Quarterly Journal of Economics 75*(3), 468–492.

Choi, H. and H. Varian (2012). Predicting the Present with Google trends. *Economic Record 88*(1), 2–9.

Frank, R. (1999). *Luxury Fever*. New York: Free Press.

Luttmer, E. F. P. (2005). Neighbors as Negatives: Relative Earnings and Well-Being. *The Quarterly Journal of Economics 120*(3), 963–1002.

Peterson, M. (2011). Working to Improve the Consumer Price Index. *Cartographic Perspectives 68*, 75–82.

Schor, J. (1993). *The Overworked American*. New York: Basic Books.

Smith, C. L. (2012). The Impact of Low-Skilled Immigration on the Youth Labor Market. *Journal of Labor Economics 30*(1), 55–89.

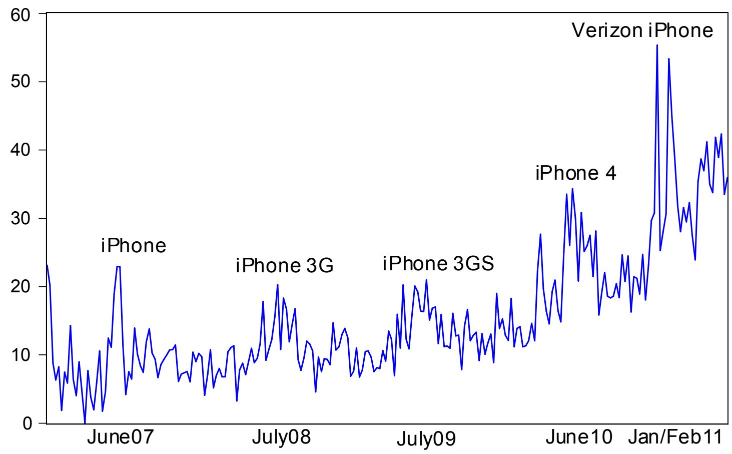

(a) Search Index for Treatment States

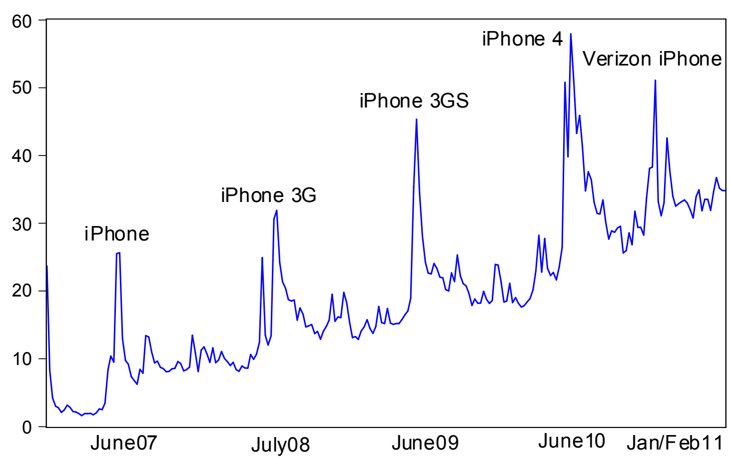

(b) Search Index for Control States

Figure 2: SEARCHES FOR "IPHONE" ACROSS TREATMENT AND CONTROL STATES : U.S., JANUARY 2007 - APRIL 2011.
Source: Google Trends

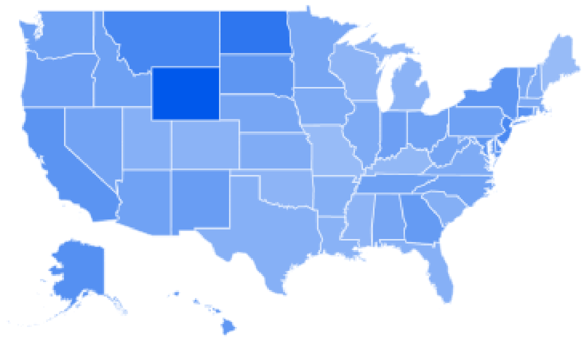

Figure 3: YOUTUBE SEARCHES FOR "VERIZON IPHONE" : JANUARY 1 - FEBRUARY 28, 2011.
Source: Google Trends

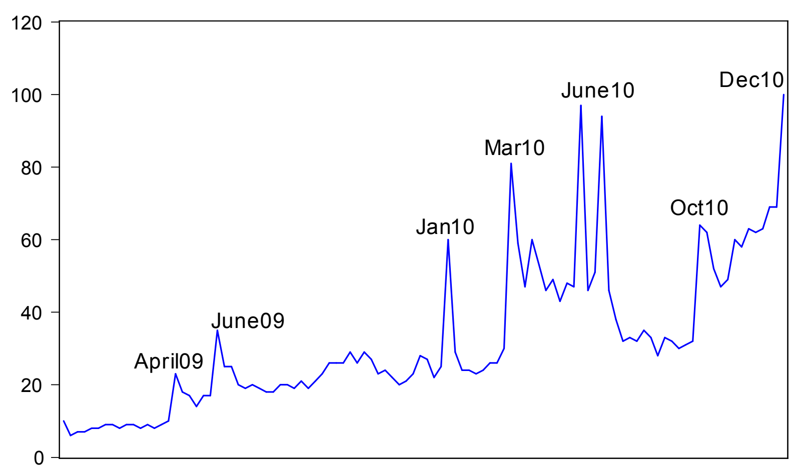

Figure 4: SEARCHES FOR "VERIZON IPHONE" : U.S., JANUARY 2009 - DECEMBER 2010.
Source: Google Trends

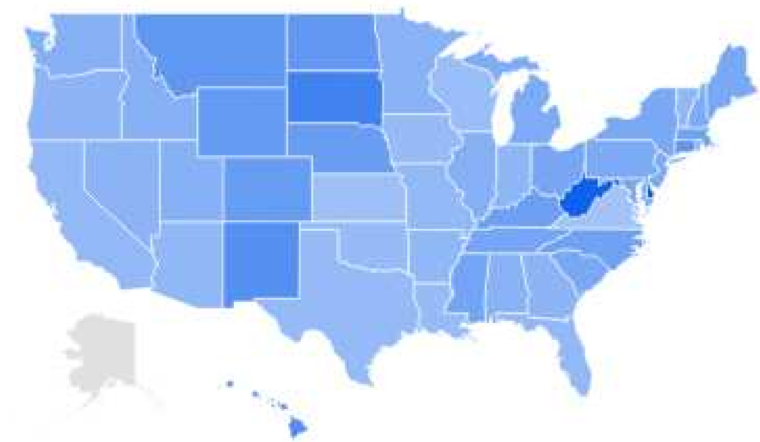

(a) Searches for “Verizon iPhone”: October 1 - December 31, 2010

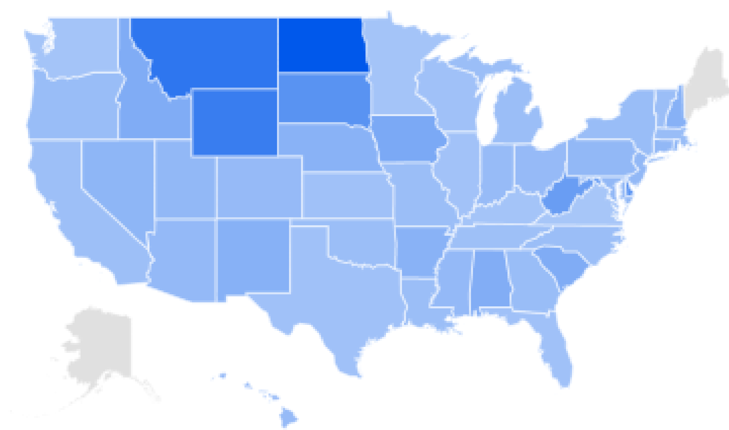

(b) Searches for “Verizon iPhone”: June 1 - July 31, 2010

Figure 5: SEARCH INTENSITIES ACROSS STATES
Source: Google Trends

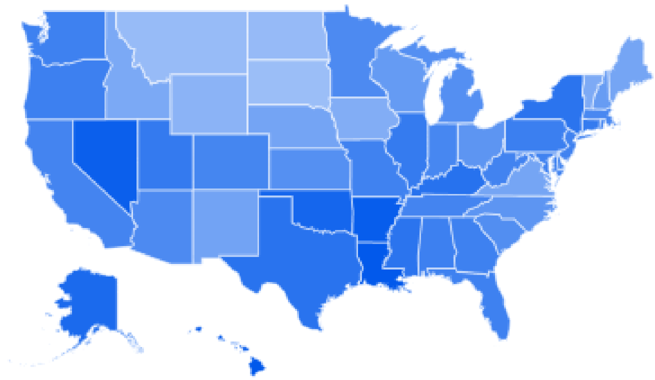

Figure 6: Searches for “iPhone 4” : June 1-July 31, 2010.
Source: Google Trends

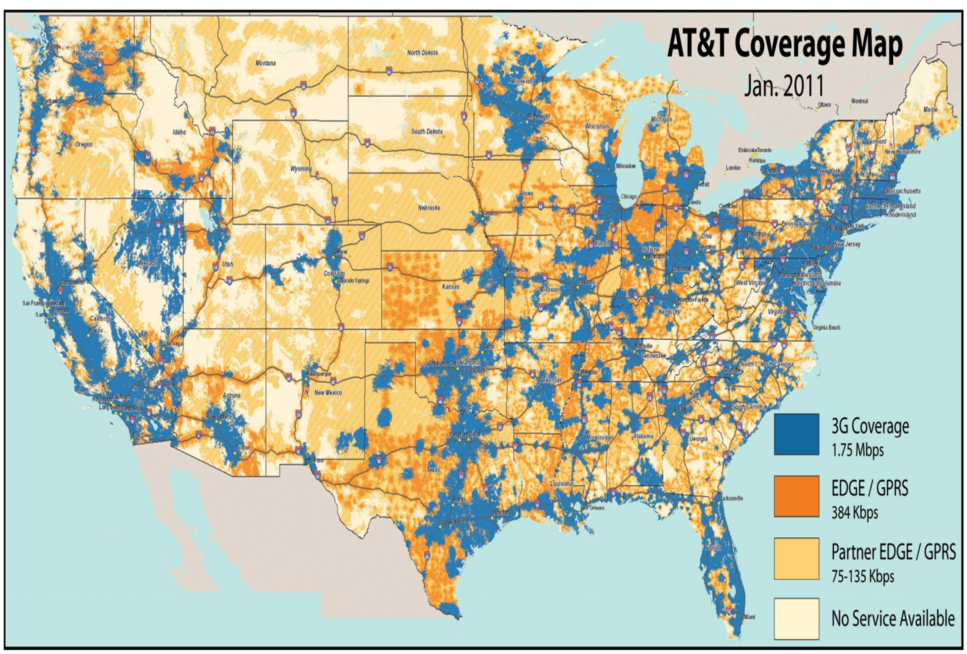

Figure 7: AT&T Coverage Map: January 2011
Source: Peterson (2011)

Table 1: Summary Statistics for Treatment and Control Groups.

This table displays personal and employment characteristics of the treatment and control groups during the treatment period. The age cohort is 15-24 and the treatment period is December 2010 - February 2011. The first four categories under the *Employment* heading represent the largest employment categories for both groups. Percentages represent proportions of each respective group.

	Treatment	Control
PERSONAL CHARACTERISTICS		
% Female	44	46
% At College	37	40
% At Primary School	21	18
% Not in Training	35	35
% Living at Home	46	57
% Head Householder	28	21
Mean Age	20	20
EMPLOYMENT		
% Eating and Drinking Places	24	29
% Educational Services	14	11
% Medical and Health Services	6	6
% Food Stores	6	7
Mean Hours	21	20
Mean Hourly Wage	\$9	\$9
Observations	627	1,581

Table 2: Treatment Effects for Different Age Groups.

This table displays the differences-in-differences coefficient β from the estimation of Eq. 3 for different age categories. The treatment intensity is 1 in January and February 2011 and .5 in December 2010. Each cell presents results from a different regression, and the baseline comprises the 15-24 age group. The dependent variable is hours per part-time worker per week and the sample period is January 1994 until December 2019. Column 7 includes all workers in the sample, aged 15-75. Column 8 presents results from a triple-differences estimation, where part-time workers aged 25-39 in the treatment states comprise an additional control group. All regressions include dummies for month, state, year, and the treatment period, together with state-year and state-month interactions and controls for the state unemployment rate and family income. Standard errors, clustered at the state level, are in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

	(1) Baseline	(2) $<$20	(3) 20-24	(4) $>$24	(5) 25-39	(6) 18-21	(7) All	(8) DDD
Treat.	2.26***	2.44**	1.70***	-0.42	-1.76***	2.84***	0.23	3.08***
	(0.39)	(1.21)	(0.28)	(0.28)	(0.51)	(0.50)	(0.25)	(0.46)
Obs.	233,361	120,030	113,331	621,845	234,119	108,578	855,281	475,834
R^2	0.04	0.05	0.04	0.01	0.01	0.04	0.01	0.07

Table 3: Treatment Effects for Different Education Categories.

This table displays the differences-in-differences coefficient β from the estimation of Eq. 3 for different education categories. Each cell presents results from a different regression. The dependent variable is hours per part-time worker per week and the sample period is January 1994 until December 2019. Columns 2 and 3 focus on those in primary and tertiary education. Column 4 excludes college students, while Column 5 applies to those not in training. All regressions include dummies for month, state, year, and the treatment period, together with state-year and state-month interactions and controls for the state unemployment rate and family income. Standard errors, clustered at the state level, are in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

	(1) Baseline	(2) Primary	(3) College	(4) Ex-College	(5) Non-Training
Treatment	2.26***	1.18	1.29*	2.38***	2.05***
	(0.39)	(0.76)	(0.73)	(0.78)	(0.58)
Observations	233,361	57,086	79,679	158,863	91,091
R^2	0.04	0.06	0.05	0.04	0.02

Table 4: Treatment Effects for Different Family Income Levels.

This table displays the differences-in-differences coefficient β from the estimation of Eq. 3 applied to different family income categories. The thresholds arise from the treatment group over the treatment period. Column 2 excludes the top 20 percent of households by family income. The treatment period is December 2010-February 2011, with respective treatment intensities of .5, 1, and 1. Each cell presents results from a different regression. The dependent variable is hours per part-time worker per week and the sample period is January 1994 until December 2019. All regressions include dummies for month, state, year, and the treatment period, together with state-year and state-month interactions and controls for the state unemployment rate and family income. Standard errors, clustered at the state level, are in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

	(1) Baseline	(2) Ex-top 20%	(3) Bottom 50%	(4) Top 50%	(5) Ex-bottom 30%
Treatment	2.26***	2.35***	1.64**	2.25***	2.52***
	(0.39)	(0.29)	(0.67)	(0.58)	(0.78)
Observations	233,361	146,037	72,365	166,177	211,009
R^2	0.04	0.03	0.04	0.04	0.04

Table 5: Treatment Effects Using Additional Control Variables.

This table displays the differences-in-differences coefficient β from the estimation of Eq. 3 using additional controls: gender, age, and hourly wage. The gender dummy equals 1 for males and 0 for females. Each cell presents results from a different regression and each column incorporates an additional control variable. The dependent variable is hours per part-time worker per week and the sample period is January 1994 until December 2019. All regressions include dummies for month, state, year, and the treatment period, together with state-year and state-month interactions and controls for the state unemployment rate and family income. Standard errors, clustered at the state level, are in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

	(1)	(2)	(3)	(4)
Additional Control:	Baseline	Gender	Age	Wage
Treatment	2.26***	2.25***	1.99***	4.32***
	(0.39)	(0.39)	(0.43)	(0.76)
Family Income	4.58***	4.74***	4.69***	4.12***
	(0.25)	(0.25)	(0.22)	(0.29)
Income-squared	-0.55***	-0.55***	-0.50***	-0.50***
	(0.03)	(0.03)	(0.02)	(0.03)
Unemployment	0.03	0.03	-0.05	0.16
	(0.07)	(0.07)	(0.05)	(0.10)
Gender		-0.28***	-0.47***	
		(0.05)	(0.05)	
Age			5.98***	
			(0.37)	
Age-squared			-0.13***	
			(0.01)	
Hourly Wage				0.44***
				(0.03)
Observations	233,361	233,361	233,361	53,549
R^2	0.04	0.04	0.14	0.07

Table 6: Robustness Tests.

This table displays the differences-in-differences coefficient from the estimation of Eq. 3. Each cell presents results from a different regression. Column 2 restricts the sample to workers who appear once during the treatment period (those in survey months 1, 3, 5, and 8.) Columns 3-6 exclude each of North Dakota, South Dakota, Montana, and Wyoming, respectively. The dependent variable is hours per part-time worker per week and the sample period is January 1994 until December 2019. All regressions include dummies for month, state, year, and the treatment period, together with state-year and state-month interactions and controls for the state unemployment rate and family income. Standard errors, clustered at the state level, are in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

	(1) Baseline	(2) Once	(3) Ex-ND	(4) Ex-SD	(5) Ex-Mont	(6) Ex-Wy
Treatment	2.26***	2.74***	2.52***	2.22***	2.00***	2.28***
	(0.39)	(0.41)	(0.33)	(0.48)	(0.36)	(0.43)
Observations	233,361	116,148	213,993	212,836	218,123	217,525
R^2	0.04	0.04	0.04	0.04	0.04	0.04

Table 7: Treatment Effects using Different Control Regions.

This table displays the differences-in-differences coefficient β from the estimation of Eq. 3 using different control regions. Each cell presents results from a different regression. Column 2 includes all areas of the control states. Column 3 uses suburban areas of all control states. Column 4 uses rural areas of all control states. Column 5 uses all areas of the six states bordering the treatment states. Column 6 uses rural areas of states with widespread AT&T coverage (Indiana, Illinois, and Michigan.) The dependent variable is hours per part-time worker per week and the sample period is January 1994 until December 2019. All regressions include dummies for month, state, year, and the treatment period, together with state-year and state-month interactions and controls for the state unemployment rate and family income. Standard errors, clustered at the state level, are in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

	(1) Baseline	(2) All Areas	(3) Suburban	(4) Rural	(5) Bordering	(6) Rural-Covered
Treatment	2.26***	1.90***	2.27***	1.46**	1.60***	2.20**
	(0.39)	(0.35)	(0.49)	(0.66)	(0.39)	(0.99)
Observations	233,361	361,098	177,611	144,618	194,729	84,305
R^2	0.04	0.04	0.05	0.04	0.04	0.04

Table 8: Treatment Effects for Different Events.

This table displays the differences-in-differences coefficient β from the estimation of Eq. 3 applied to different events. The first event is the introduction of the Verizon iPhone, where the treatment period is December 2010 to February 2011, with respective treatment intensities of .5, 1, and 1. The second event is the introduction of the iPhone 4, where the treatment period is June and July 2010, with each month subject to a treatment intensity of 1. Columns 1 and 2 presents results from the events in isolation. For both events, the treatment group comprises the Dakotas, Montana, and Wyoming. The dependent variable is hours per part-time worker per week and the sample period is January 1994 until December 2019. All regressions include dummies for month, state, year, and the treatment period, together with state-year and state-month interactions and controls for the state unemployment rate and family income. Standard errors, clustered at the state level, are in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

	(1)	(2)	(3)
Treatment Months	Verizon iPhone	iPhone 4	Both
Dec10-Feb11	2.30***		2.26***
	(0.38)		(0.39)
Jun10-July10		-1.23***	-1.14***
		(0.39)	(0.39)
Observations	233,361	233,361	233,361
R^2	0.04	0.04	0.04

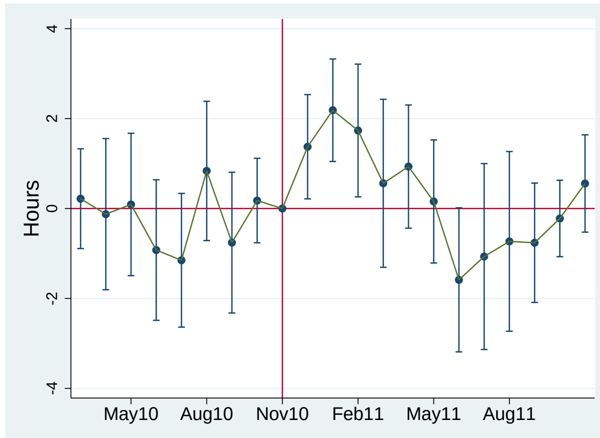

(a) Treatment Effects for Part-Time Workers aged 15-24

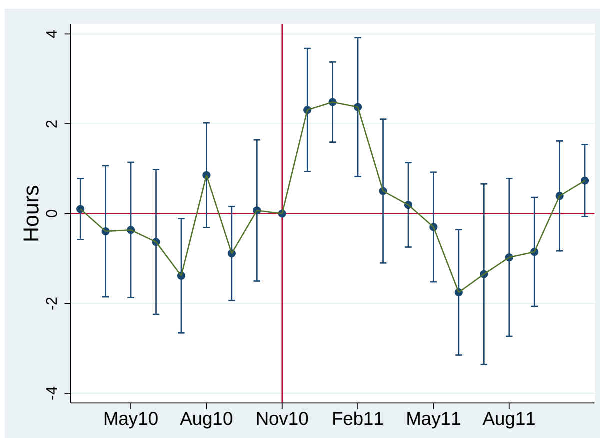

(b) Treatment Effects Excluding Bottom 30% of Workers by Household Income

Figure 8: Estimated monthly treatment effects for 15-14 cohort, together with 95 percent confidence intervals. Each point estimate represents the weekly hours difference per part-time worker between treatment and control groups in that month relative to the same difference in November 2010. Panel b) presents results when excluding the bottom 30 percent of households by family income. Treatment effects are estimated from Eq. 4; see Section 3.4.1 for more details. The first bar represents the period prior to April 2010, while the last represents the period after October 2011. Controls consist of the state unemployment rate, income, and income squared. All regressions include dummies for month, state, year, the treatment period, the treatment group, together with state-month interactions and state-specific yearly trends. Standard errors are clustered at the state level.

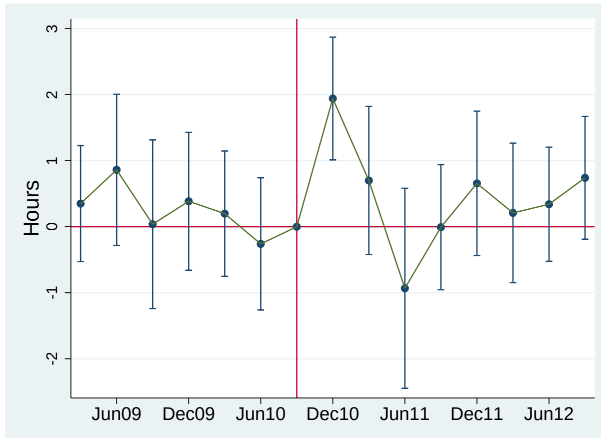

(a) Treatment Effects for Part-Time Workers aged 15-24

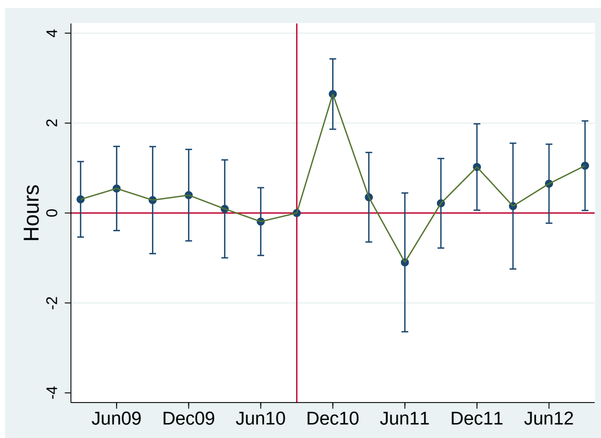

(b) Treatment Effects Excluding Bottom 30% of Workers by Household Income

Figure 9: Estimated quarterly treatment effects for 15-24 cohort, together with 95 percent confidence intervals. Each point estimate represents the average weekly hours difference between the treatment and control groups during the three-month period commencing that month relative to the same difference per part-time worker for the three-month period commencing September 2010. Treatment effects are estimated from Eq. 4; see Section 3.4.1 for estimation details. Panel b) presents results when excluding the bottom 30 percent of households by family income. The first bar represents the period prior to June 2009, while the last represents the period after August 2012. Controls consist of the state unemployment rate, family income, and family income squared. All regressions include dummies for month, state, year, and the treatment period, together with state-month interactions and state-specific yearly trends. Standard errors are clustered at the state level.

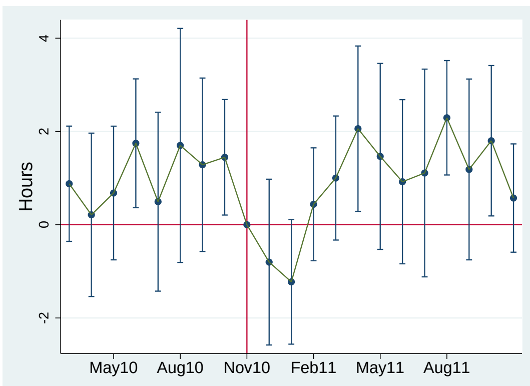

(a) Treatment Group is the 25-39 Cohort in Treatment States and Control Group is the 25-39 Cohort in Control States.

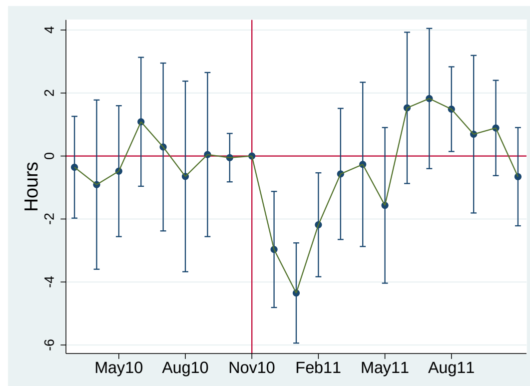

(b) Treatment Group is the 25-39 Cohort and the Control Group is 15-24 Cohort, both from Treatment States

Figure 10: Estimated monthly treatment effects for 25-39 cohort, together with 95 percent confidence intervals, using different control groups. Each point estimate represents the weekly hours difference between treatment and control groups in that month relative to the same difference in November 2010. Treatment effects are estimated from Eq. 4; see Section 3.4.1 for estimation details. Panel a) displays the treatment effects when the treatment and control groups comprise the 25-39 cohort. Panel b) presents results when the control group is the 15-24 cohort *within* the treatment group. The first bar represents the period prior to April 2010, while the last represents the period after October 2011. Controls consist of the state unemployment rate, family income, and family income squared. All regressions include dummies for month, state, year, and the treatment period, together with state-month interactions and state-specific yearly trends. Standard errors are clustered at the state level.

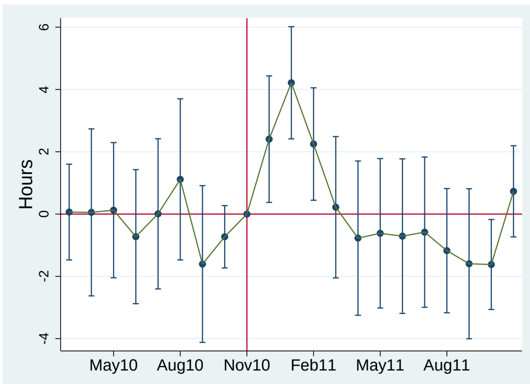

(a) Triple Difference Treatment Effects for Workers aged 15-24

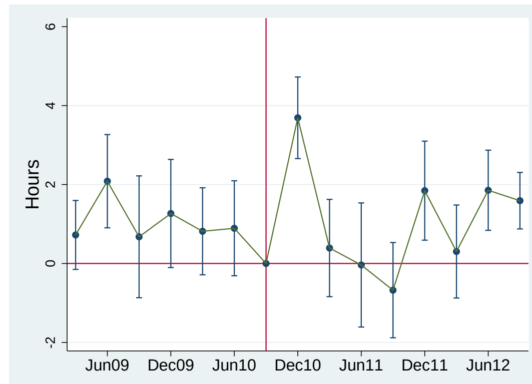

(b) Triple Difference Treatment Effects Excluding Bottom 30% of Workers by Household Income

Figure 11: Treatment effects from a triple difference estimation, together with 95 percent confidence intervals. The control group comprises the 15-24 cohort in the control states and the 25-39 cohort in the treatment states. Each estimate gives the hours difference between those in the 15-24 cohort in the treatment and control states, less the hours difference for the 25-39 group in the treatment states between the treatment and control periods. Panel a) presents results when estimation occurs at monthly frequency. The first bar represents the period prior to April 2010, while the last represents the period after October 2011. The baseline is November 2010. Panel b) presents results when estimation occurs for quarterly intervals. The first bar represents the period prior to June 2009, while the last represents the period after August 2012. The baseline is the three-month period starting September 2010. Controls consist of the state unemployment rate, income, and income squared. All regressions include dummies for month, state, year, and the treatment period, together with state-month interactions and state-specific yearly trends. The dotted lines indicate the 95 percent confidence intervals. Standard errors are clustered at the state level.