Leisure Luxuries and the Labor Supply of Young Men*

Mark Aguiar Mark Bils Kerwin Kofi Charles Erik Hurst

February 22, 2018

Abstract

Since 2004, time-use data show that younger men, ages 21 to 30, shifted their leisure sharply to video gaming and other recreational computer activities. Over the same period, these younger men exhibited a larger decline in work hours than older men or women. We propose a framework to answer whether improved leisure technology affected younger men's labor supply. The starting point is a leisure demand system that parallels that often estimated for consumer expenditures. We show that total leisure demand is especially sensitive to innovations in leisure luxuries, that is, activities that display a disproportionate response to changes in total leisure time. Using cross-region variation, we estimate that gaming/recreational computer use is distinctly a leisure luxury for younger men. We calculate that innovations to gaming/recreational computing reduced younger men's labor supply by 1.5 to 3.1 percent since 2004. That would explain 23 to 46 percent of their decline in market hours over the period.

*We thank Shirley Yarin and Hyun Yeol Kim for outstanding research assistance. We are grateful to Bob Hall for detailed and constructive comments on an earlier draft. We also thank Thomas Crossley, Matt Gentzkow, Patrick Kehoe, John Kennan, Pete Klenow, Alan Krueger, Hamish Low, Kevin Murphy, and Yona Rubinstein, as well as seminar participants at Berkeley, Board of Governors of the Federal Reserve, Boston University, Chicago, Columbia, Harvard, Houston, IIES Stockholm, LSE, Penn, Princeton, Stanford, UCL, UIC, Wharton, and the Federal Reserve Banks of Atlanta, Chicago and Richmond for helpful comments. Authors contact information: maguiar@princeton.edu, mark.bils@gmail.com, kerwin.charles@gmail.com, and erik.hurst@chicagobooth.edu.

1 Introduction

We ask if innovations to leisure technology, specifically to recreational computing and gaming, reduced the labor supply of younger men. Our focus is propelled by the sharp changes we see in time use for younger men, ages 21 to 30, during the 2000s. Comparing the American Time Use Surveys (ATUS) for 2012 to 2015 to eight years prior, 2004 to 2007, we see for younger men that: (a) hours of market work decreased by 2.5 hours per week,[1] or 6.7 percent, more than twice the decline for men ages 31 to 55;[2] (b) the decline in market hours was mirrored by a nearly equivalent increase in leisure; and (c) increased time spent in gaming and computer leisure comprised 80 percent of that increase in leisure. Younger men increased their time spent in gaming and computer leisure by 46% percent over this short period. Non-employed young men now average 520 hours a year in recreational computer time, sixty percent of that spent playing video games. This exceeds their time spent on home production or socializing with friends. Older men and women allocate much less time to computer leisure and displayed much less growth in these activities during the 2000s.

An elemental question is whether increased computer use and gaming reduced younger men's market hours, or simply reflected their response to working fewer hours due, say, to reduced labor demand. That is, has improved leisure technology raised the return to non-market time, thereby increasing the reservation wages of younger men, or are we witnessing movement along a stable labor supply curve? To identify shifts in labor supply we introduce a leisure demand system that parallels that typically considered for consumption expenditures. We show that total leisure demand is especially sensitive to innovations to technology for leisure luxuries. Leisure luxuries are activities that exhibit little diminishing returns to time and therefore display disproportionate responses to changes in total leisure time.

We estimate how various leisure activities respond to total leisure time, tracing out "leisure Engel curves." Our estimates exploit region-industry-year variations in leisure, such as that caused by differential impact of the Great Recession across US states and sectors. The identifying assumption is that such cross state-industry variation in total leisure is not driven by differential changes in preferences or technologies across leisure activities. We find that gaming and recreational computer use is distinctively a leisure luxury for younger men, but only a modest luxury for other groups. A one percent increase in leisure time is associated with more than a 2 percent increase in time spent playing video games for younger

[1] This secular decline in market hours for younger men is sizable by historical standards. For comparison, consider the double-dip recession in the early 1980s: between 1979 and 1982, men ages 21-55 in the Current Population Surveys (CPS) decreased their hours worked by approximately 3.2 hours per week on average.

[2] Women ages 21-30 and ages 31 to 55 show much smaller declines of 1.2 and 1.7 percent, respectively. Data from the CPS, described below, also show considerably larger declines in hours for men ages 21 to 30 than seen for older men or for women.

men, roughly double the elasticity found for other demographic groups.

We next divide the large increase in recreational computer use by younger men into a movement along a leisure Engel curve versus a shift in that curve. We judge the shift in technology for computer leisure relative to that for leisure devoted to sleeping and personal care, a technology we assume is fairly static. Our leisure Engel curves predict that computer recreation would have increased by 8 percent between 2004 and 2015 in response to younger men's total increase in leisure. Thus the bulk of the actual increase, 38 out of 46 percent, is attributed to better technology for computer leisure.

The last step is to map the impact of better leisure technology on the value of total leisure and, thereby, on labor supply. The mapping from the return to leisure to labor supply depends on how changes in market work affect younger men's consumption. We consider two scenarios. If individuals are "hand-to-mouth," so consumption equals labor earnings, we calculate that improvements in computer leisure since 2004, holding wages fixed, predict a 1.5 percent decline in the market hours of younger men. That translates to 23 percent of the decline in market work observed for younger men in the ATUS. Alternatively, if the marginal utility of consumption is held constant, which in our framework holds a dollar's marginal value constant, then the impact is twice as large, yielding a 3.2 percent decline in market work for younger men, which translates to 46 percent of their decline in market work. So we conclude that better leisure technology was a significant factor, though not necessarily the primary factor, in the decline in hours for younger men. By contrast, we find that better computer technology had no effect on the labor supply of older men and only a small effect on that of women, results compatible with our finding that the activity is not a strong leisure luxury for either group. Consequently, improvements in recreational computing technology can account for the majority of the *differential* decline in the labor hours of younger men relative to older men.

While we focus on the impact of computer leisure on younger men since the early 2000's, our approach should be more broadly applicable. For instance, if the relevant time-use data were available, one could estimate changes in the return to leisure stemming from prior leisure innovations such as the introduction of television. How these prior leisure innovations translated to changes in observed labor market outcome depended not only on how the growth in leisure technology affected labor supply but also on how other contemporaneous forces were affecting labor demand.

A natural question is how younger men support themselves in the face of declining earnings. We document that 67 percent of non-employed younger men lived with a parent or close relative in 2015, compared to 46 percent in 2000. We also examine expenditures for households with younger men in the Panel Study of Income Dynamics. We see little, if any, decline in the relative consumption of younger men since 2000. This lends some credence to the assumption that younger men's consumption has been little affected by their increased gaming. We complement the PSIS results with reported life satisfaction information from the General Social Survey. Despite stagnant wages, declining employment rates, and an increased propensity to live with their parents, younger men report increased happiness during the 2000s. This contrasts sharply with older men, whose satisfaction fell along with their relative earnings. We see the life satisfaction results as indirect evidence that younger men experienced relatively little decline in the consumption and greatly valued their improved leisure options.

Our focus on time allocation owes a natural debt to the seminal papers of Mincer (1962) and Becker (1965). They emphasize that labor supply is influenced by how time is allocated outside of market work—for instance, female labor force participation being affected by improved household technology. Our work complements that of Greenwood and Vandenbroucke (2008), Vandenbroucke (2009), and Kopecky (2011), who use a quantitative Beckerian model to show that declines in relative prices of leisure goods help to explain declining employment over the last century. We add to this literature by introducing and estimating a leisure demand system, showing that labor supply is most affected by technology for leisure luxuries, and illustrating how one can relate shifts in labor supply to changes in the allocation of time across leisure activities.

The paper is organized as follows: Section 2 examines changes in time use during the 2000s, emphasizing the sharp increase in computer and gaming time for younger men; Section 3 presents our methodology including the leisure demand system; Section 4 highlights our identification strategy and estimates the leisure Engel curves; Section 5 uses the demand system and changes in time allocation to infer changes in leisure technology; Section 5 also quantifies the shifts in leisure and labor supply curves for different demographic groups during the 2000s; Section 6 highlights the robustness of our results to alternate parameterizations; Section 7 examines cohabitation, consumption, and self-reported well being for younger men; and Section 8 concludes.

2 Younger Men's Changing Composition of Leisure

We first document how younger men, and other demographic groups, have allocated their non-market time since the early 2000's based on the time diaries of the American Time Use Survey (ATUS) from 2004 through 2015. The ATUS draws a sample from CPS respondents and surveys them within a few months after the final CPS survey, collecting a 24-hour diary in which respondents record the previous day's activities in 15-minute intervals. The ATUS groups these activities into categories.[3] We restrict the sample to civilians ages 21 to 55. We further exclude full-time students who are less than age 25.[4] This mitigates any role for increased college attendance in the decline in work hours for younger men.

2.1 Trends in Broad Time Use Categories

We divide activities into six broad categories: market work, job search, home production, child care, education, and leisure. Job search includes sending out resumes, job interviewing, and researching jobs. Home production includes doing household chores or maintenance, preparing meals, shopping, and caring for other adults.[5] We separate child care from home production. Education refers to time spent on one's *own* education, such as attending courses or doing homework. Leisure consists of watching television and movies, recreational computing and video games, reading, playing sports, hobbies, etc. We discuss leisure in more detail in the next subsection. We treat a portion of eating, sleeping, and personal care (ESP) as leisure as these categories have both a biological and leisure component. To isolate the leisure component of ESP, we exclude 7 hours per day from total ESP time to account for the fact that a certain amount of sleeping, eating and personal is needed for survival.[6]

Table 1 shows time use for younger and older men (Panel a) and younger and older women (Panel b). We report time use in weekly hours, multiplying the daily averages by 7. To increase power we group data for 2004-2007 and 2012-2015. The table reports average time for each category by time period, as well as differences across the two periods.

Starting with the top panel, we see that younger men reduced their market work by 2.5 weekly hours over this period, which corresponds to a nearly 7 percent decline. Comparing top and bottom rows of the panel, we see this decline in market hours was nearly matched by an increase in leisure of 2.3 hours for younger men.[7] The remaining time activities display relatively small changes. By comparison, older men reduced their weekly market work by 1.1 hours, while increasing their leisure by 1.2 hours. Panel (b) shows patterns for women. Younger women had a smaller decline in market work, but a larger decline in home

[3] Each ATUS sample is uniformly distributed across days of week. Time spent traveling to or from an activity is always included in the activity's time. Though the ATUS starts in 2003, we begin our analysis with 2004, as there are small changes in the survey methodology between 2003 and 2004.

[4] Before 2013 the CPS, and therefore the ATUS, asked only those under age 25 about school attendance. The Data Appendix discusses in more detail our ATUS sample, as well as other data sets employed in the paper. Throughout the paper, we weight observations by the relevant survey's sampling weight.

[5] Some small categories like personal health care and unclassified time use are omitted from our analysis.

[6] Approximately 95 percent of respondents report 7 or more hours per day for ESP. We explored alternative adjustments (e.g., excluding 6 or 8 hours per day for biological ESP needs) and found our results were not sensitive to these changes.

[7] Appendix Figure A1 displays the cross-sectional distribution of leisure time for younger men for the 2004-07 and 2012-15 sub-periods. The density displays a noticeable rightward shift over time.

Table 1: Broad Time Allocation During the 2000s, Hours Per Week

(a) Men, Age 21-55

	Age 21-30			Age 31-55		
Activity	2004-2007	2012-2015	Change	2004-2007	2012-2015	Change
Market Work	38.4	36.0	-2.5	40.9	39.7	-1.1
Job Search	0.3	0.7	0.4	0.3	0.4	0.1
Home Production	12.1	11.4	-0.7	14.8	13.9	-0.9
Child Care	2.8	2.4	-0.4	3.6	4.1	0.4
Education	2.5	3.2	0.7	0.6	0.6	0.0
Leisure	61.0	63.4	2.3	57.0	58.1	1.2

(b) Women, Age 21-55

	Age 21-30			Age 31-55		
Activity	2004-2007	2012-2015	Change	2004-2007	2012-2015	Change
Market Work	27.4	27.1	-0.3	27.4	27.0	-0.5
Job Search	0.2	0.3	0.1	0.2	0.3	0.1
Home Production	19.0	17.5	-1.5	24.2	22.4	-1.8
Child Care	10.0	8.8	-1.1	7.4	7.6	0.2
Education	2.3	2.9	0.6	1.1	1.0	-0.1
Leisure	58.5	59.9	1.4	56.1	58.0	1.9

Note: Table reports hours per week spent on activities from the ATUS. Data are pooled for 2004-2007 and 2012-2015 periods. The difference between the two periods is the reported change. An individual's total time endowment, after subtracting off 49 hours per week for biological sleeping, eating and personal care needs, is 119 hours per week.

production, than younger men. The decline in home production was even more pronounced for older women, generating a larger increase in leisure than for younger women or older men. Comparing across all groups, younger men exhibited the largest gain in leisure.

To explore the robustness of the trends in market work by differing demographic groups, we use data on annual hours worked from the March CPS. An advantage of the March CPS over the ATUS is that market hours are reported based on the calendar year, rather than a snapshot from a single day.[8] Comparing years 2012-2015 to 2004-2007, the March CPS show a decline 179 hours per year (3.4 hours per week) for men ages 21-30, which is a somewhat steeper decline than seen in the ATUS. As with the ATUS data, the decline in market hours for the younger men is considerably larger than 105 hours per year (2 hours per week) decline observed for men ages 31 to 55. The relative difference in the decline in market hours between young and older men was nearly identical in the ATUS and March CPS (1.2 hours per week vs. 1.4 hours per week). The March CPS data also allows us to break the trends in hours worked between trends in employment versus hours per worker. The employment rate fell by 5.2 percentage points for younger men from 2004-2007 to 2012-2015; and constituted 60 percent of their decline in annual hours. For older men employment declined by 3.1 percentage points.

From the March CPS responses on weeks worked during the prior year, one can also see the prevalence of longer-term non-employment. Figure 1 plots the fraction of younger and older men who worked *zero* weeks over the year. About 8 percent of each age group report zero weeks worked in 2000. The share not working increased considerably for both groups during the 2000s; but the increase is much more dramatic for younger men. For younger men that fraction began increasing prior to the Great Recession, accelerated during the Great Recession, and has only modestly recovered. As of 2015, the share of younger men not working the entire year was nearly 15 percent.

2.2 Trends in the Nature of Leisure

We now explore leisure at a more disaggregated activity level. Within total leisure, we distinguish the following five activities: recreational computer time; television and movie watching; socializing; discretionary eating, sleeping and personal care (ESP); and other leisure. Recre-

[8]Our measure of annual hours worked in the March CPS is the respondent's report of their usual hours per week worked multiplied by the number of weeks they worked during the prior calendar year. As with the ATUS sample, we exclude full-time students ages less than 25 when using the March CPS sample. In the appendix we discuss the trends in hours in the CPS in greater detail. We also also report hours trends from the American Community Surveys (ACS). The ACS data allow us to show robustness to excluding all full-time students from the sample, not just those ages less than 25. The trend in hours are very similar across the various samples.

Figure 1: Fraction of Men With Zero Weeks Worked Over Prior Year by Age, March CPS

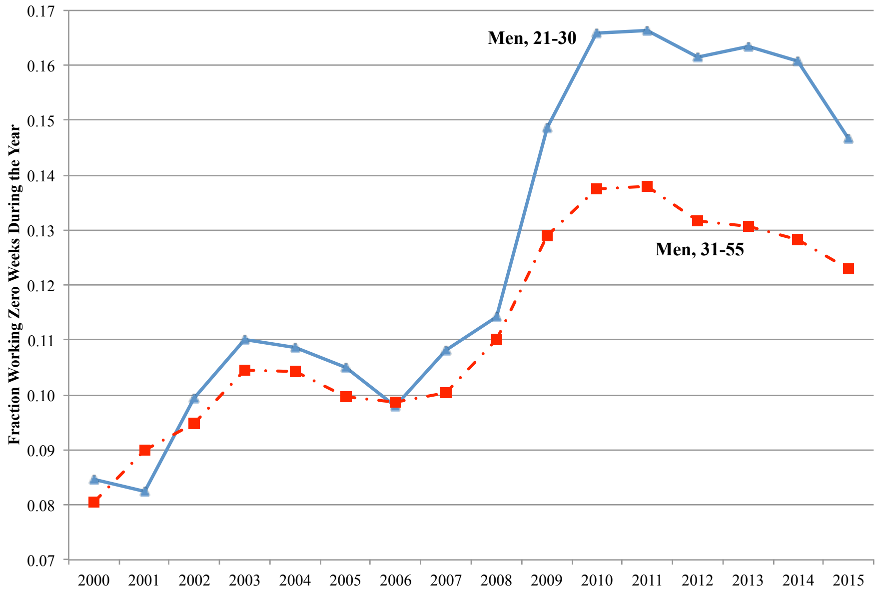

Note: The figure shows the shares of men ages 31-55 (squares) and men ages 21-30 (triangles) who report working zero weeks during the prior year. Data are from the CPS March supplement. Full-time students ages less than 25 are excluded.

ational computer time includes time spent on non-work email, playing computer games, browsing web sites, leisure time on smart phones, online chatting, and engaging in social media. We often highlight the video/computer game component of recreational computer.[9] Computer time for work or non-leisure activities (like paying bills or checking email) are captured by other time-use categories. Watching television and movies includes watching streaming platforms, as well as traditional television and movies. Socializing includes entertaining or visiting friends and family, parties, dating, and participating in civic or religious activities. "Other leisure" includes all remaining leisure activities, such as reading, listening to music, exercising, playing sports, and engaging in hobbies.

Table 2 shows the weekly hours spent by younger men in each leisure category. We see that the increase of 2.3 hours in weekly leisure hours for younger men is largely accounted for by an increase of 2.0 hours in their recreational computer time. This represents over 80 percent of the total leisure increase for younger men.[10] Furthermore, most of that increase

[9] The ATUS has a category of time use labeled "playing games." This includes video games, but also includes playing cards as well as traditional board games like checkers, Scrabble, etc. So we cannot distinguish playing Scrabble from video gaming. We document a very large increase in playing games during the 2000s by younger men. We equate this with an increase in video gaming. However, we realize that we may be identifying a Scrabble boom as opposed to a video game boom.

[10] Appendix Figure A2 displays the cross-sectional distribution of recreational computing time for younger men, conditional on spending a strictly positive amount of time. Similar to the leisure distribution presented in Figure A1, the distribution of computing time displays a prominent rightward shift between 2004-07 and

took the form of increased video game playing (1.4 hours per week). The implied annual increase in computer leisure of 104 hours is a striking change for a time-use category over a short span of time. For reference, annual hours women spend at home production fell by 520 hours over the last forty years (Aguiar and Hurst (2007)). As a corollary of that increase in computer time, other leisure categories changed very little despite the large increase in total leisure. Younger men did not spend more time watching TV/movies, socializing, or at other leisure activities. The only other leisure category that recorded a notable increase is eating, sleeping, and personal care; but that increase, 0.5 hours, represents only a two percent increase over the sample period.

Why did recreational computing display such explosive growth for younger men over this period? One major innovation in the mid 2000s was people moving their social interactions, and especially their gaming, online. Facebook, started in 2004, grew from 12 million users in 2006 to 360 million by 2009. A generation of video game consoles introduced in 2005 and 2006 allowed individuals to interact online. Massive multiplayer online games launched around the same time. For example, World of Warcraft began in 2004 and grew to 10 million monthly subscribers by 2010. Coupled with advances in graphics, these innovations fueled a large expansion of the video game industry. Nominal revenues of the video game industry increased by about 50 percent between 2006 and 2009 after being fairly flat for the prior five years.[11]

Table 2: Leisure Activities for Men 21-30, Hours per Week

Activity	2004-2007	2012-2015	Change
Total Leisure	61.0	63.4	2.3
Recreational Computer	3.3	5.2	2.0
Video Game	2.0	3.4	1.4
ESP	24.3	24.9	0.5
TV/Movies/Netflix	17.3	17.1	-0.2
Socializing	7.8	7.9	0.1
Other Leisure	8.3	8.2	-0.1

Note: Leisure components sum to total leisure time. Video gaming is a subcomponent of total computer time. ESP refers to eating, sleeping and personal care net of 49 hours.

From Table 2, weekly leisure hours for younger men increased by 2.3 hours between 2004-07 and 2012-15. At the same time, there was a large increase, from 10.3 to 14.0 percent, in

2012-15.

[11]Data are from the NPD group: vgsales.wikia.com/wiki/NDP_sales_figures.

the share of younger men in the ATUS who are not employed. Because the non-employed exhibited nearly 30 hours more leisure on average in 2004-2007, the shift to fewer employed played a major role in the overall increase in average leisure. In Table 3 we look at leisure conditional on employment status. Unfortunately, since there is no panel dimensions to the ATUS, we are comparing different pools of employed and non-employed individuals across a period with a large decrease in employment. So it is important to keep in mind that the changes in average leisure calculated for those employed and not employed will reflect compositional effects driven by the greater share not employed. This is especially true for statistics calculated for the non-employed, as this group expanded by nearly 40 percent. To the extent this growth came from younger men with relatively strong attachments to the labor force, and to associated leisure choices, shifting composition will act to substantially decrease average leisure calculated conditional on non-employment in the latter period.

Turning to Table 3, we see that leisure for employed younger men increased by 2.0 weekly hours, 65 percent of which is accounted for by increased recreational computing. As anticipated, the non-employed have substantially more leisure. *Conditional on non-employment*, leisure hours actually fell since 2004. However, as discussed above, this may reflect a composition shift in the pool of non-employed, as its share dramatically grew. For example, if the additional non-employed exhibit 25 percent less leisure than the balance of that group, that alone would be sufficient to reduce the conditional average by about 6.2 hours per week, which is greater than the measured decrease. That 25 percent differential, while just an example, represents the difference between the 25th percentile of the non-employed leisure distribution for 2004-07 and the average for that sub-sample.

Consistent with this, looking at the last row of Table 3, the non-employed in 2012-2015 were much more likely to allocate time to education and job search, with most of this reflecting education. These increases exactly offset the decline in leisure time. We suspect this increase in education reflects a combination of the changing composition of those not employed and a response to perceived benefits of schooling.[12] Despite the overall decline in leisure time for non-employed younger men during the 2000s, time spent on recreational computers (video games) increased for this group by 4.2 (2.4) hours per week. It is also worth noting that in 2012-2015 non-employed young men spent nearly 10 hours per week (520 hours per year) on recreational computer activities. This exceeds both the amount of time they spend socializing on non-computer activities and the amount of time they spend on other leisure categories.[13]

[12] Note, the latter may provide a further factor to explain declining market hours for younger men. We do not see this as conflicting with an impact of better leisure technology, especially as we will attribute only a portion of the decrease in younger men's market hours to better leisure options.

[13] That average time spent on computer leisure by non-working younger men masks a great deal of het-

Table 3: Leisure Activities for Men 21-30 (Hours per Week): By Employment Status

Activity	Employed 2004-2007	Employed 2012-2015	Employed Change	Non-Employed 2004-2007	Non-Employed 2012-2015	Non-Employed Change
Total Leisure	57.6	59.6	2.0	87.0	82.1	-4.9
Recreational Computer	3.0	4.3	1.3	5.4	9.6	4.2
Video Game	1.8	2.9	1.0	3.5	5.9	2.4
ESP	23.6	23.9	0.3	30.2	29.9	-0.2
TV/Movies/Netflix	15.9	15.5	-0.4	27.8	25.0	-2.8
Socializing	7.4	7.8	0.3	10.6	8.9	-1.7
Other Leisure	7.7	8.1	0.5	13.0	8.6	-4.4
Job Search and Education	2.0	1.9	-0.1	9.2	14.1	4.9

Note: Components sum to total leisure time. Video gaming is a subcomponent of total computer time. ESP refers to eating, sleeping and personal care net of 49 hours per week.

Below we infer changes in computer leisure technology from how individuals shifted leisure toward that activity, adjusting for changes in total leisure time. As a first look at the data, we sort individuals into bins based on hours of leisure their previous day. The bins are on the horizontal axis of Figure 2, where, for example, label 5 indicates individuals who spent five to six hours at leisure. For each leisure bin, we report average time spent at recreational computer use. The lighter bars in the figure depict the averages for younger men for 2004-2007, while the darker bars depict those for 2012-2015. We see that computer leisure increased within essentially all leisure bins, but especially for high-leisure individuals.

Table 4 compares younger men's shift toward computing and gaming (top panel) to that for older men, younger women, and older women (bottom three panels). The table clearly shows that the increase in computer leisure in general, and its gaming component in particular, was a younger men's phenomenon. While younger men increased their computer leisure by 2.0 hours per week, the increases were only 0.1, 0.6, and 0.5 hours per week for older men, younger women, and older women, respectively. Women reported a modest increase in their recreational computer time; but, in contrast to younger men, zero of that increase involved video games.

erogeneity. In 2004-2007, 30 percent of non-working younger men reported computer leisure for the prior day; for 2012-2015 that number is 40 percent. Conditional on spending such time, non-working younger men spent 2.6 and 3.4 hours for the day, respectively, in 2004-2007 and in 2012-2015. During the 2012-2015 period, 11 percent of non-working younger men spent more than 4 hours for the day at computer leisure.

Table 4: Computer Leisure and Video Game By Age-Sex-Skill Groups, ATUS

	2004-2007	2012-2015	Change
	Men 21-30		
Total Leisure	61.0	63.4	2.3
Recreational Computer	3.3	5.2	2.0
Video Games	2.0	3.4	1.4
	Men 31-55		
Total Leisure	57.0	58.1	1.2
Total Recreational Computer	2.1	2.2	0.1
Video Games	0.9	0.8	-0.1
	Women 21-30		
Total Leisure	58.5	59.9	1.4
Total Recreational Computer	1.5	2.2	0.6
Video Games	0.8	0.8	0.0
	Women 31-55		
Total Leisure	56.1	58.0	1.9
Total Recreational Computer	1.6	2.1	0.5
Video Games	0.6	0.7	0.1

Note: Video game time is a subcomponent of computer leisure.

Figure 2: Younger Men's Hours per Day of Computer Leisure by level of Total Leisure

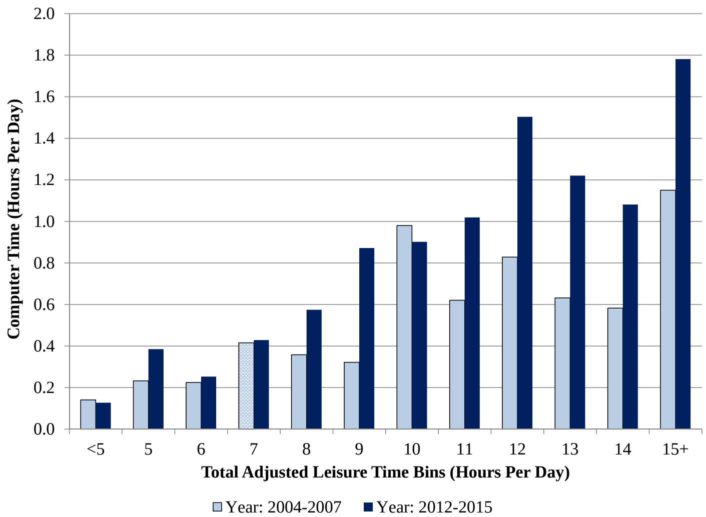

Note: Figure shows average time spent on computer leisure (including video games) by individual's total leisure. Time use is expressed in hours per day. Except for first and last bins, leisure bins span one hour per day, with minimal value of each bin denoted.

3 Leisure Luxuries and Labor Supply

In this section we derive a leisure demand system that maps total leisure into specific leisure activities. We show how observations on changing time allocations can be used to infer shifts in the quality of leisure activities and, in turn, changes in the marginal return to total leisure. The change in that marginal return can then be linked to shifts in labor supply. This section develops the theoretical groundwork for the empirical estimation in Section 4 and the quantitative results of Sections 5 and 6. Given that our data reflects each individual's activities for a single day, we consider a static framework. Dynamic factors such as habit formation are natural theoretical extensions, but less amenable to empirical analysis given the nature of our dataset.

3.1 Preferences

Agents have preferences over a consumption good, c, and time spent on leisure activities h_i, $i = 1, ..., I$. We assume separability between consumption and leisure activities, writing utility as $U(c, v(h_1, ..., h_I; \boldsymbol{\theta}))$ where v is an aggregator over leisure activities and $\boldsymbol{\theta} = \{\theta_1, ..., \theta_I\}$ is a vector of technology shifters. While not necessary for all results, for simplicity we impose strong separability between c and v by setting $U_{cv} = 0$.

We assume v has the following functional form:

$$v(h_1, ..., h_I; \boldsymbol{\theta}) = \sum_{i=1}^{I} \frac{(\theta_i h_i)^{1-\frac{1}{\eta_i}}}{1 - \frac{1}{\eta_i}}. \tag{1}$$

The parameter $\eta_i > 0$ is activity specific and governs the diminishing returns associated with additional time spent on activity i. Increases in the technology parameter θ_i increase the utility associated with spending a given amount of time on activity i.

While each leisure activity enters with its specific elasticity η_i, the activities are assumed to be additively separable from one another (although the entire v function may be raised to a power, which would be a feature of the overall utility function U). This assumption implies that the marginal value of allocating time to one leisure activity over another is not dependent on how leisure time is allocated across the remaining activities. We provide some empirical support for this assumption in Section 4.

3.2 Leisure Engel Curves

The agent faces a wage w in terms of the consumption good, and chooses how to allocate their time endowment (normalized to one) across the I leisure goods and the labor market. If N denotes market labor, the time constraint is $N + \sum_i h_i = 1$. We assume $N > 0$ is optimal and omit $N \geq 0$ as a constraint. As discussed below, however, much of the analysis carries over to an environment in which market labor is fixed or rationed. The functional form for v requires $h_i > 0$ at an optimum.

Specifically, the agent's problem is:

$$\max_{c, h_1, ..., h_I} U(c, v(h_1, ..., h_I; \boldsymbol{\theta})) \tag{2}$$

subject to

$$c + \sum_{i=1}^{I} p_i \leq w\left(1 - \sum_{i=1}^{I} h_i\right) + y, \tag{3}$$

where y is initial wealth or non-labor income and p_i is the price of technology bundle θ_i. For the present, we treat $\boldsymbol{\theta}$ as parameters that must be purchased, but will later discuss the choice of $\boldsymbol{\theta}$. The problem as stated can be interpreted as the optimal allocation problem conditional on a vector $\boldsymbol{\theta}$, with a subsequent step of optimizing over the possible technology bundles.

Let λ denote the multiplier on the budget constraint. The first-order conditions are:

$$U_c = \lambda \tag{4}$$

$$U_v v_i = \lambda w \text{ for } i = 1, ..., I, \tag{5}$$

where $v_i = \partial v / \partial h_i$, and similarly $U_c = \partial U / \partial c$ and $U_v = \partial U / \partial v$.

In the spirit of Browning et al. (1985), it is useful to analyze this problem in stages. In particular, let $\omega \equiv \lambda w$ denote the opportunity cost of time, which is the the marginal value of wealth times the wage. Given this price of time, the agent makes a labor-leisure decision subject to $H \equiv \sum_i h_i = 1 - N$. Given H, the agent allocates leisure time across individual activities, equating the marginal utilities.

More formally, let μ denote U_v at the optimal allocation, which is the marginal return to increasing the leisure aggregate v. Rewriting (5), we have:

$$h_i = \theta_i^{\eta_i - 1} \mu^{\eta_i} \omega^{-\eta_i}. \tag{6}$$

For a given μ, the elasticity of demand for activity i with respect to the shadow price ω is η_i. Activities with relative high η_i are the ones that are most sensitive to the opportunity cost of time. All else equal, an increase in technology θ_i increases or decreases time allocated to the associated activity depending on whether $\eta_i \gtrless 1$. If a leisure activity becomes more enjoyable, whether one spends more or less time at that activity turns on the size of the elasticity, with one being the crucial threshold.

Summing over the various leisure activities, (6) implies:

$$H = \sum_i \theta_i^{\eta_i - 1} \mu^{\eta_i} \omega^{-\eta_i}. \tag{7}$$

By the envelope condition of the leisure allocation sub-problem, μ is pinned down by total leisure H, which in turn is determined by ω. Equation (7) implicitly defines H as a function of ω given $\boldsymbol{\theta}$. Thus, we can write $H(\omega, \boldsymbol{\theta})$ as the optimal choice of leisure given the shadow price of time ω and technology $\boldsymbol{\theta}$. Similarly, let $h_i(\omega, \boldsymbol{\theta})$ denote the demand for activity i given by (6).

We can use equation (6) to trace out a "leisure Engel curve." Consider individuals with different ω but employing the same leisure technology vector $\boldsymbol{\theta}$. That is, individuals differ by wages or wealth that shift the shadow price of time. From (6) and (7), one obtains:

$$\frac{\partial \ln h_i}{\partial \ln \omega} = \frac{\eta_i}{\overline{\eta}} \frac{\partial \ln H}{\partial \ln \omega}, \tag{8}$$

where $\overline{\eta} \equiv \sum_j \eta_j s_j$ is the weighted average elasticity, and the weights are given by the share of leisure time devoted to each activity: $s_j \equiv h_j/H$.

Equation (8) will play an important role in our empirical work. As we look across agents with different values of time, we observe how time allocated across individual leisure activities varies with total leisure time. This elasticity is the activity's own price elasticity divided by the weighted average of all elasticities. Activities with a greater η_i increase disproportionately with total leisure. That is, high η_i activities are "leisure luxuries." Our notion of a leisure luxury parallels that of a consumption luxury (or superior) good in traditional consumption demand systems. Given its importance, we denote this elasticity by β_i:

$$\beta_i \equiv \frac{\eta_i}{\overline{\eta}}. \tag{9}$$

Our derivation of the leisure Engel curves does not hinge on how total hours of leisure H are determined. If labor time were indivisible or rationed, we would treat N as a parameter of the individual's problem. This would imply the constraint $\sum_i h_i \leq 1 - N$ be added to problem (2). Let ω be the multiplier on that constraint and the analysis goes through unchanged. The crucial assumption is that the shadow price of time is the same when choosing between alternative leisure activities, not whether the price of time is pinned down by the wage, labor market frictions, or the returns to home production.

3.3 Inferring Technological Progress

The agent's time allocation problem also sheds light on technological progress in leisure activities. Let I denote the activity of interest, which in the empirical analysis will be recreational computer use. Let $j \neq I$ be a "reference activity." In the empirical implementation, we consider several alternatives as the reference. From the respective first-order conditions (6):

$$\frac{\ln h_I}{\eta_I} - \frac{\ln h_j}{\eta_j} = \left(\frac{\eta_I - 1}{\eta_I}\right) \ln \theta_I - \left(\frac{\eta_j - 1}{\eta_j}\right) \ln \theta_j. \tag{10}$$

The fact that the common price of time, ω, and the marginal utility of v, μ, are differenced out implies that this equation holds independently of wages, non-labor income, and the levels of consumption and leisure. It exploits the fact that the returns to individual activities are equated at the margin.

Now consider how time allocation changes as technology changes. Differencing (10) gives:

$$\frac{\Delta \ln h_I}{\eta_I} - \frac{\Delta \ln h_j}{\eta_j} = \left(\frac{\eta_I - 1}{\eta_I}\right) \Delta \ln \theta_I - \left(\frac{\eta_j - 1}{\eta_j}\right) \Delta \ln \theta_j. \tag{11}$$

The left-hand side is the change in relative time allocation between activity I and the reference activity j, normalized by the elasticities. The right-hand side captures the change in relative technologies.

Equations (8) and (11) play an important role in our empirical analysis. To gain intuition for how technology can be inferred from time allocation, consider shocks to the price of time ω, such as due to job loss or consumption changes. The shift in relative time allocation between activities I and j will be determined by the relative Engel curve elasticities β_I and β_j. These are movements *along* the leisure Engel curves. Equation (11) shows that the change in how time is allocated to activity I relative to j, in *excess* of that explained by the relative slopes of their respective Engel curves, reflects changes in relative technology. These are *shifts* in the Engel curves rather than movements along the curves. A goal of the empirical exercise is to separate movements along leisure Engel curves due to changes in leisure hours from shifts in technology.

An alternative approach to inferring technological change uses prices. Consider changes over time in θ_i, and the decision of the agent to purchase the latest technology. In particular, suppose the agent can upgrade technology by $\Delta\theta_i$ by paying an additional Δp_i. The utility gain from a marginal improvement is $U_v \frac{\partial v_i}{\partial \theta_i} \Delta\theta_i$. The opportunity cost is $\lambda \Delta p_i$, where λ is the shadow value of wealth. Using the first-order conditions, the agent prefers the marginal upgrade as long as:

$$\frac{\Delta\theta_i}{\theta_i} \geq \left(\frac{p_i}{wh_i}\right) \frac{\Delta p_i}{p_i}. \tag{12}$$

The term in parentheses reflects relative cost shares in producing the leisure activity; that is, the numerator is the cost of the technology and the denominator is the cost of the time input, priced at the market wage. For the marginal consumer, equation (12) will hold with equality. Equation (12) provides an alternative to measuring technological change. It does not exploit the time allocation decision, but uses the ability to substitute between time inputs and market inputs in the production of leisure. We explore both approaches in the empirical work.

3.4 The Response of Labor Supply to Leisure Technology

The derivation of the Engel curve elasticities took the relative shadow price of time ω as a parameter and traced out the choice over individual leisure activities. We now return to the problem of choosing consumption and market hours.

Recall from (7) that the choice of leisure is pinned down by the opportunity cost of time $\omega = \lambda w$ and the technology vector $\boldsymbol{\theta}$. Let ϵ denote the absolute value of the Frisch elasticity of leisure $\epsilon \equiv -d \ln H / d \ln \omega$, keeping in mind that this elasticity holds constant the marginal value of income, λ. Implicitly differentiating (7) and using $\mu = U_v$, we have:

$$\frac{1}{\epsilon} = \frac{1}{\overline{\eta}} - \frac{\partial \ln U_v}{\partial \ln v} \frac{\partial \ln v}{\partial \ln H}. \tag{13}$$

The left-hand side is the inverse Frisch elasticity of leisure. The first term on the right-hand side is the inverse elasticity of the leisure aggregate v with respect to H. The second term captures the elasticity of the marginal utility of leisure with respect to H. The role of $\overline{\eta}$ comes directly from the leisure demand system. If the agent devotes a high share of leisure time to leisure luxuries then, all else equal, $\overline{\eta}$ is larger and so is the Frisch elasticity. This reflects that leisure luxuries are elastic at the margin, and a small change in the opportunity cost of time induces a large shift in the amount of leisure time.[14]

The second term on the right side of (13) shows that the response of leisure depends not only on the curvature of the leisure aggregate v (the first term, $1/\overline{\eta}$) but also on how v enters the utility function $U(c, v)$. If, for example, $U(c, v) = u(c) + v$, then the final term on the right is zero and the Frisch elasticity of leisure is governed solely by $\overline{\eta}$. If v enters U via a concave function, then $\epsilon < \overline{\eta}$, as the additional concavity mitigates the response of leisure to a change in wage. Conversely if v enters U via a convex function, then $\epsilon > \overline{\eta}$, although the extent of convexity must be bounded such that the sum on the right is positive.

To trace out the impact of a change in θ_i on H, we differentiate (7) and, using (13), rearrange to obtain:

$$\frac{\partial \ln H(\omega, \boldsymbol{\theta})}{\partial \ln \theta_i} = s_i \left(\epsilon \beta_i - 1 \right). \tag{14}$$

The response of H to a change in technology depends on the share of that activity in leisure, s_i; the Frisch elasticity, ϵ; and the elasticity of the Engel curve, $\beta_i = \frac{\eta_i}{\overline{\eta}}$.

[14] There are close antecedents to this result in the literature on consumption. In particular, Crossley and Low (2011) discuss the restrictions necessary for a constant elasticity of inter-temporal substitution of expenditures in a demand system involving multiple consumption goods. Browning and Crossley (2000) demonstrate the link between relative income elasticities and willingness to substitute inter-temporally. Both points have clear parallels in our treatment of labor supply with multiple leisure goods.

The role of the share s_i is intuitive. If the activity is only a small component of overall leisure, then improving its technology has a minimal impact on total leisure.

The Frisch elasticity enters for the same reason it governs the response to a wage change. The change in leisure technology shifts the relative return of leisure versus work, and the Frisch captures the sensitivity of leisure to that relative price. Note that the same force is relevant if we were to model labor as an extensive-margin choice. For example, if we consider a distribution of agents with heterogeneous wages and wealth, an increase in the return to leisure will sweep the marginal agents out of employment. The Frisch elasticity is then determined by the density of agents with reservations wages close to the market wage, as in Chang and Kim (2006).

Finally, the elasticity of the Engel curve, β_i, plays an important role in the response of leisure to technological changes. This elasticity captures the extent of diminishing returns to that activity relative to the overall elasticity. An activity with a high η_i does not experience strong diminishing returns. Thus, total leisure must increase significantly in order to restore equality between the marginal return to leisure and the opportunity cost of time. The presence of β_i indicates that the response of total leisure to technological change is particularly strong for improvements in a leisure luxury.

From equation (14), we can derive the response of labor supply to a change in technology. The overall response depends on how consumption responds as well as leisure. At one extreme, we hold the marginal utility of consumption constant. This implies that in the consumer's problem, non-labor income adjusts as well; that is, the individual is perfectly insured against changes in technology. To be precise, y adjusts to external changes in technology, and not in response to the agent's labor-leisure decision. For this "full insurance" scenario, the partial derivative $\partial H(\omega, \boldsymbol{\theta})/\partial \theta_i$ appropriately captures the response of leisure. From equation (14):

$$\Delta \ln H_{FI} \approx s_i \left(\epsilon \beta_i - 1\right) \Delta \ln \theta_i, \tag{15}$$

where the FI notation indicates that the change assumes full insurance such that the marginal utility of consumption is held fixed.

Absent insurance, we need to consider the full derivative of $H(\omega, \boldsymbol{\theta})$, including the impact on the marginal utility of consumption, which is embedded in ω. Including the wealth effect will dampen the response of leisure. To explore this channel, we consider the opposite extreme to full insurance by letting consumption move one-for-one with labor earnings. We refer to this scenario as "hand-to-mouth" as movements in earnings are reflected fully in consumption. In particular, suppose $\Delta c = -w\Delta H$.[15] Letting ρ denote the inter-temporal elasticity of consumption, $\rho \equiv -U_c/(U_{cc}c)$, and differentiating (7) yields:

$$\Delta \ln H_{HTM} \approx \frac{\Delta \ln H_{FI}}{1 + \frac{\epsilon}{\rho}\left(\frac{H}{1-H}\right)\left(\frac{w(1-H)}{c}\right)}. \tag{16}$$

Thus, relative to full insurance, the "hand-to-mouth" sensitivity of leisure to θ is scaled down by the wealth effect, which depends on the ratio of the curvature parameters ρ and ϵ, as well as the ratio of leisure to work and the ratio of labor income to consumption.

Combining (11) with (15), we obtain:

$$\Delta \ln H_{FI} = s_I \left[\frac{\epsilon\beta_I - 1}{\overline{\eta}\beta_I - 1}\right]\left(\Delta \ln h_I - \frac{\beta_I}{\beta_j}\Delta \ln h_j\right). \tag{17}$$

The alternative measure of technological progress (12) yields the following:

$$\Delta \ln H_{FI} \approx s_I\left(\epsilon\beta_I - 1\right)\left(\frac{p_I}{wh_I}\right)\Delta \ln p_I, \tag{18}$$

where $\Delta \ln p_I$ is the change in price across vintages of θ_I. The hand-to-mouth calculations are scaled down by dividing by the denominator of (16).

Comparing (17) and (18), we see that former requires a measure of $\overline{\eta}$, the average elasticity, while the latter requires a measure of the relative cost shares and the additional cost of new technology. We shall explore both approaches in the empirical analysis.

The framework presented in this section provides an empirical road map. In the next section, we take the leisure demand system of Section 3.2 to the data to estimate β_i for the leisure activities discussed in Section 3. In Section 5 we use equation (11) and the empirical shift in time allocation to estimate the change in technology for recreational computer use and video games. We combine this with price data and use (12) to recover $\overline{\eta}$. The last step is to use (15) and (16) to quantify the impact of improved technology on labor supply.

4 Estimating Leisure Engel Curves

We now estimate the leisure demand system outlined in Section 3.2. The key targets are the Engel curve elasticities β_i. From estimates of the Engel curves, we will construct estimates of the primitives θ_i and η_i. In this section, we discuss in turn measurement error, functional

[15] More generally, $\Delta c = -w\Delta H - \Delta p_I$, where Δp_I is the change in cost due to the upgrade in technology. Including this effect involves subtracting $\Delta p_I/(\rho c)$ from the numerator of (16). This adjustment is likely to be small as expenditure c is much larger than the marginal price change Δp_I of new vintages.

forms, and identification. We then report our estimated Engel curve elasticities.

4.1 Measurement Error

The major measurement challenge is that the time diaries are a single-day's snapshot, with zeros reported for most activities on that given day. Ideally, we would like data on an individual's typical allocation of leisure, which requires observations over multiple days or even weeks. The lack of such broader coverage makes our data especially prone to sampling error. A secondary concern is that measurement error in an individual activity will distort measured total leisure as well, given that total leisure is simply the sum of the individual activities. This generates an artificial correlation, a well known issue in estimating consumption demand systems.

To address both issues, we construct synthetic time diaries that average over similar types of individuals. Specifically, we form cells based on gender, age, educational attainment, industry, geographic region, and time period. Age is demarcated as in Section 2.1; namely, 21-30 and 31-55. Educational attainment is split by those with at least a bachelor's degree versus those with less than 16 years of schooling, omitting full time students throughout. Industry is reported as of the last CPS interview, typically a few months prior to the time diary. The CPS asks the industry of the current job or, if not currently employed, the industry of the last job held in the preceding 12 months.[16] Note that we include as a separate "industry" a missing industry code, which typically reflects those who have not had a job in the preceding 12 months.

For region, we first compute the change in each state's average leisure between 2004/07 and 2012/15 separately for each gender-age group. We then sort states into five roughly equally-populated groups based on the recorded change. Thus individuals in states with a large increase in leisure are grouped separately than those in states with a small increase (or decrease) in leisure.

The final cell characteristic is time period, where we use the three periods discussed in Section 2.1; namely, 2004-07, 2008-11, and 2012-15.

Theoretically, this implies up to 1,680 cells; but in practice, some cells contain no individuals. In estimating, we weight all cells by the sum of its individual members' weights and restrict attention to cells with at least 10 observations.

[16]Specifically, we use PRMJIND1 in the ATUS-CPS file. The 13 industries are: 1.) Agriculture, forestry, fishing, and hunting; 2.) Mining; 3.) Construction; 4.) Manufacturing; 5.) Wholesale and retail trade; 6.) Transportation and utilities; 7.) Information; 8.) Financial activities; 9.) Professional and business services; 10.) Educational and health services; 11.) Leisure and hospitality; 12.) Other services; and 13.) Public administration. The final CPS industry is Armed Forces, which is not present in our sample. We treat individuals without an industry code as the 14th industry.

4.2 Specification

Our empirical specification builds on the consumption literature, most notably Deaton and Muellbauer's (1980) Almost Ideal Demand System (AIDS). Adapting AIDS to a leisure demand system, we posit that the share of time allocated to an activity is approximately linear in the log of total leisure time. Letting i denote an activity, t a time period, and k demographic cell, we estimate:

$$s_{ikt} = \delta_{it} + \gamma_i \ln H_{kt} + \varepsilon_{ikt}, \tag{19}$$

where $s_{ikt} = h_{ikt}/H_{kt}$ is the share of total leisure H_{kt} devoted to activity i in period t by group k, while $\ln H_{kt}$ is log of that group's total leisure time. We include time-period fixed effects, δ_{it} in all specifications. As added controls, we consider fixed effects for education, region, and industry, respectively, across alternative specifications. Time-dependent shifters that influence the allocation of leisure time to activity i are captured by δ_{it}. In particular, δ_{it} controls for movements in technology θ_i. We estimate (19) separately for each activity and allow all parameters to vary by age-gender groups.

From estimate $\hat{\gamma}_i$, we recover an estimate of $\beta_i = \partial \ln h_i / \partial \ln H$:

$$\hat{\beta}_i = 1 + \frac{\hat{\gamma}_i}{\bar{s}_i}, \tag{20}$$

where $\bar{s}_i$ is the average of activity i's leisure share over the sample period, specific to each age-gender group.

4.3 Identification

To consistently estimate γ_i from (19) requires that H_{kt} is orthogonal to the error term. Recall that the activity-time fixed effect δ_{it} captures time-dependent shifts in tastes or technology that are uniform across cells.[17] Thus, our identifying assumption is that cell-specific tastes for a given leisure activity are uncorrelated with total leisure.

To flesh out our identification assumption, note that an ideal source of variation in a cell's relative leisure time would be forces such as differential employment opportunities due, say, to the Great Recession. This type of variation allows an accurate measure of how leisure

[17] In the case of computers and video games, the assumption of common technology seems justified, given the widespread and rapid diffusion of these technologies during the 2000s. According to the FCC, all MSAs had high speed internet as of 2000. We explored using regional variation in introducing broadband internet as a shift in the quality of recreational computing. However, since broadband had saturated the country by the start of our time use data, that leaves no regional or time-series variation to use as an instrument.

is allocated across activities due to exogenous changes in total leisure, where by exogenous we mean independent of idiosyncratic tastes and technologies for a particular activity.

The construction of our cells is designed to isolate such variation. In particular, the 2000s saw large relative swings in employment across education groups, regions, and industries. These movements are plausibly unrelated to idiosyncratic shifts in the taste for particular leisure activities.[18] Thus, by grouping individuals in cells defined by education, industry, and region, not only minimizes measurement error, but also isolates a plausibly exogenous source of variation in total leisure.

The threat to identification arises if cells with especially high total leisure systematically have different tastes and technologies for an activity than cells with low levels of leisure. For example, suppose that cells with high leisure have a relative preference for recreational computing. In this case, we will over-estimate the Engel curve elasticity for computing and under-estimate the elasticities for other activities. Conversely, if high-leisure cells have a weaker taste for computing, we will under-estimate the Engel elasticity for that activity, and over-estimate the other activities' elasticities. To the extent our cells are broadly defined and designed to isolate variation due to aggregate labor market conditions, such a failure of orthogonality should not be a primary concern. To address concerns that the level of leisure may be correlated with demographic characteristics, we explore the robustness of the results to adding fixed effects for education, industry, and region. With these controls, the concern for orthogonality arises only if a differential correlation still remains after controlling for the average level of the leisure activity within that education, industry, or regional group.

4.4 Estimates

Table 5 reports our estimates of β_i for younger men for each of the leisure activities reported in Table 2. We also break out video gaming from its broader computer category. All estimates are based on the AIDS specification, equation (19), and the implied $\hat{\beta}_i$ are obtained using (20).[19] The first column is a baseline specification that includes time-period fixed effects. The second column adds education-group fixed effects; the third column further adds regional fixed effects; the final column adds fixed effects for the fourteen industry groups. Thus, by the final column, all variation is based on time series variation within the subgroups relative to the average cell effect and the aggregate time fixed effect. The standard errors for β_i are

[18] This assumption is supported by evidence suggesting that much of the cross-state variation in market work during the 2000s was driven by industrial composition or housing markets. See, for example, Charles et al. (2016) and Mian and Sufi (2014).

[19] Estimates of the AIDS coefficients, γ_i, are reported in Appendix Table A1. We also estimated (19) allowing γ_I to vary across time. An F-test that the coefficient is the same across the three time-periods has a p-value of 0.40.

bootstrapped.[20]

Table 5: Leisure Engel Curves of Younger Men: $\hat{\beta}_i$

	(1)	(2)	(3)	(4)
Recreational Computer	2.46	2.46	2.24	1.64
	(0.17)	(0.17)	(0.19)	(0.21)
Video Games	2.90	2.82	2.42	2.03
	(0.21)	(0.22)	(0.22)	(0.29)
TV/Movies/Netflix	1.18	1.08	1.15	1.14
	(0.06)	(0.06)	(0.06)	(0.06)
Socializing	0.72	0.75	0.52	1.00
	(0.11)	(0.11)	(0.12)	(0.14)
ESP	0.70	0.72	0.78	0.76
	(0.05)	(0.05)	(0.05)	(0.07)
Other Leisure	0.98	1.11	1.12	1.05
	(0.08)	(0.08)	(0.09)	(0.10)
Fixed Effects:				
Time Period	✓	✓	✓	✓
Education		✓	✓	✓
Geographic			✓	✓
Industry				✓
Number of Cells	242	242	242	242
Number of Individuals	6,250	6,250	6,250	6,250

Note: Implied $\hat{\beta}_i$ using AIDS specification. An observation is a time-gender-age-education-industry-state group cell. Bootstrapped standard errors are in parentheses.

As seen from Table 5, computers and video games are leisure luxuries. Focusing on the results in Column 1, recreational computing has an Engel elasticity of 2.46, while the video games sub-component has an elasticity of 2.90. The estimates suggest that video game time is the most luxurious leisure activity for younger men. All other activities have elasticities close to or strictly less than 1. TV/Movie watching has an estimated leisure elasticity of 1.18. Other Leisure is neither a luxury nor necessity ($\hat{\beta}_i = 0.98$). Eating-sleeping-personal care is a leisure necessity ($\hat{\beta}_i = 0.70$), as is socializing ($\hat{\beta}_i = 0.72$).

The Engel curve elasticities are similar across specifications, save perhaps for the last column. Including industry fixed effects moves the estimated elasticities towards one. Most

[20] Specifically, the bootstrap procedure repeatedly draws samples, estimates the AIDS coefficient γ_i and the average share $\bar{s}_i$, and computes $\hat{\beta}_i$ using equation (20). The bootstrap is performed using the 160 replication weights provided by the ATUS.

of this movement occurs by including a dummy for those without an industry code; that is, individuals who have been non-employed for at least twelve months. The fact that these individuals have disproportionately high leisure and devote relatively more time to computing implies that including their fixed effect "flattens" the estimated Engel curve. In terms of the calculation of equation (17), the shallower slope for recreational computing, relative to that for ESP, implies that less of the observed increase in recreational computing should be attributed to moving "along" its Engel curve, with more attributed to improvements in its technology. In this sense, the estimates of Column 1 are more conservative than Column 4 for estimating the impact of this better leisure technology on labor supply.[21]

Table 6 reports the estimated Engel elasticities of computing and ESP for other demographic groups. The specification is that of Column 1 from Table 5. The implied elasticity for recreational computing is 1.37 for older men, 1.22 for younger women, and 1.41 for older women, all of which are roughly half that estimated for younger men. ESP is a leisure necessity for all groups.

Table 6: Engel Curve Estimates by Demographic Group

	Men 31-55	Women 21-30	Women 31-55
Recreational Computer	1.37	1.22	1.41
	(0.10)	(0.22)	(0.08)
ESP	0.59	0.71	0.63
	(0.02)	(0.03)	(0.02)
Number of Cells	388	220	354
Number of Individuals	32,614	9,357	39,427

Note: Specification is that of Table 5 Column 1. Bootstrapped standard errors are in parentheses.

Figure 3 provides a visual sense of the data behind the estimation of the computer Engel curve for younger men. Specifically, it depicts a scatter plot of log recreational computer time against log total leisure. Each point represents a cell average. Circles depict $2004-2007$ observations; triangles depict those for $2012-2015$. The two fitted lines imply estimated elasticities of 2.44 and 2.53 respectively for the earlier and later periods. A test that the slopes are different has a p-value of 0.89; so the hypothesis that time allocated to recreational computer shifted up proportionally across states cannot be rejected. This figure clearly shows

[21] If we use Column (4)'s estimates, the implied shift in labor supply is -3.31%, compared to the -3.06% reported in Table 7 Column (2).

Figure 3: Leisure Engel Curves for Computer Leisure: 2004-2007 vs. 2012-2015

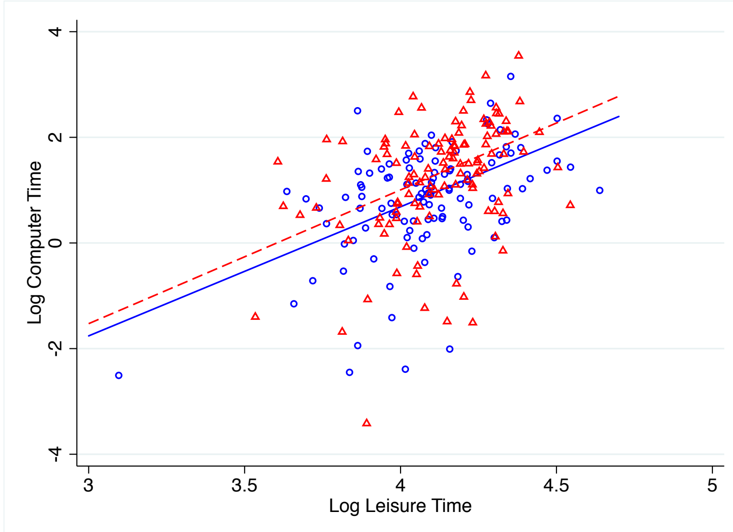

Note: Figure depicts a scatter plot of cell average leisure time (horizontal axis) and recreational computing and gaming (vertical axis), both in log hours per week. All cells are included regardless of size. The circles represent data from 2004-2007 and the triangles represent 2012-2015. The solid line is the weighted regression line for the earlier period and the dashed for the later period. The slopes with standard errors are 2.44 (0.40) and 2.53 (0.50), respectively. A test of whether the slopes are equal has a p-value of 0.89.

a shift upwards of the computing Engel curve for young men during the 2000s. These patterns provide a sense of how we disentangle movements along a stable Engel curve from shifts due to increasing θ_I.

As a further check on our leisure demand system, we re-visit the assumption of additive separability across activity sub-utilities (equation 1). This implies that, conditional on H, time spent at activity i offers no information on the relative returns to activities j versus k ($j, k \neq i$). To explore if this is consistent with the data, we ask if time spent at computer leisure predicts how remaining leisure is divided across other activities. Specifically, we group the younger men, combining years 2004 to 2015 of the ATUS, into three groups based on computer leisure (h_I) the prior day: $h_I = 0$, $h_I \in (0, 2]$ hours per day, and $h_I > 2$ hours/day. Denote these groups by $n = 0, 1, 2$, respectively. The first group comprises roughly 70 percent of the sample, while the latter two each comprise about 15 percent. For each group we compute $h_{in}/(H_n - h_{In})$ for i = TV/movies, socializing, ESP, and other leisure–that is, shares of non-computer leisure time devoted to that activity.[22] Figure 4 reports the mean shares for each group. The figure controls for the fact that groups with greater computer use, because they also have greater total leisure, should allocate more of remaining leisure

[22] h_{in} is the average time spent on activity i for group n; H_n is the average leisure time for group n; and h_{In} is average computer time for group n.

Figure 4: How Non-Computer Leisure is Allocated to Other Categories, Younger Men

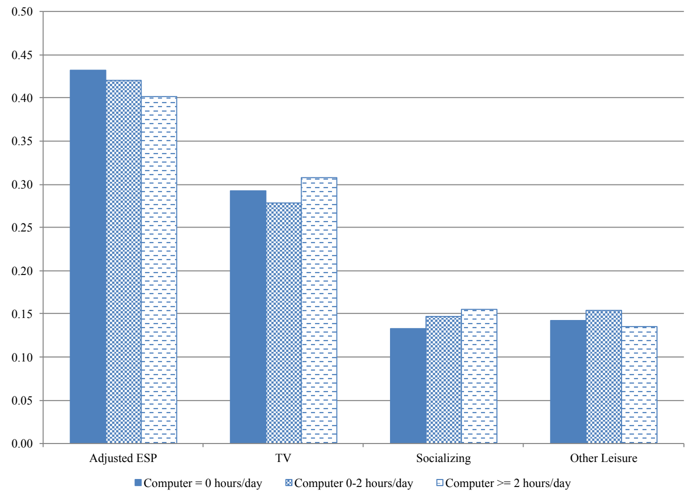

Note: Data pool the 2004-2015 ATUS. The sample excludes full-time students, ages less than 25. We stratify by three groups: younger men who spent zero time on computer leisure the prior day, those who spent 2 hours or less, and those who spent more than 2 hours. Time allocated at an activity is adjusted for group differences in total leisure as described in footnote 23.

to watching TV/Movies and less to ESP, given our estimated leisure Engel curves.[23] There is little systematic differences in how non-computer leisure is allocated by those who spend no time, some time, and a great deal of time at computer leisure. This is consistent with our assumption of separability between computer and other leisure activities.

5 Leisure Luxuries and Labor Supply During the 2000s

In this section, we use time diaries and the leisure demand system developed and estimated above to infer technological progress for computer leisure. We then assess the impact of this change on labor supply.

[23]Specifically, we estimate the AIDS specification by regressing each activity's share of non-computer leisure $(h_{ikt}/(H_{kt} - h_{Ikt}))$ on total non-computer leisure time $(\ln[H_{kt} - h_{Ikt}])$. Let $\hat{b}_i$ denote the estimated coefficients. Then $\hat{b}_i\,[\ln(H_n - h_{In}) - \ln(H_0 - h_{I0}]$ is the predicted difference in shares based on the estimated Engel curves. We subtract this from the shares yielding the results in the figure. Note that if differences in time allocation line up on the Engel curves, the three columns will be the same height for each activity.

5.1 Implied Technological Change from Time Use

With the estimates of $\hat{\beta}_i$ in hand, we can use time-series trends in time allocation to infer the rate of technological progress for gaming and computer leisure since the early 2000s. We begin with equation (11), which relates changes in time allocation to changes in technology. As noted in Section 3, changes in time allocations identify relative technology changes. For our baseline, we treat leisure eating/sleeping/personal care (ESP) as our reference activity. This assumes no technological or preference change for eating, sleeping or personal care during our sample period. Setting $\Delta\theta_{ESP} = 0$ in (11) and indicating activity I as recreational computer, we have:

$$(\eta_I - 1)\Delta \ln \theta_I = \Delta \ln h_I - \frac{\beta_I}{\beta_{ESP}} \Delta \ln h_{ESP}. \tag{21}$$

As reported in Table 2, younger men increased ESP time by 2.2 percent over the ATUS sample period. The estimates in Table 5 imply that $\hat{\beta}_I/\hat{\beta}_{ESP} = 3.5$. This implies that, absent any technological change, their computer time would increase by 7.8 percent. This is the final term on the right-hand side of equation (21), and corresponds to the predicted movement along the Engel curve for computer leisure. However, computer time for younger men actually rose by 46.6 percent. We therefore estimate the change in $(\eta_I - 1)\Delta \ln \theta_I$ to be 38.7 percent (with standard error of 7.1 percent), or 4.8 percent per year.[24,25]

We can repeat this calculation for other demographic groups. For example, we estimate for younger women that $(\eta_I - 1)\Delta \ln \theta_I$ increased by 23.4 percent (standard error 6.0 percent), or 3.0 percent per year. The only group which does not show an increase in computer technology is older men. For this group, $(\eta_I - 1)\Delta \ln \theta_I = -6.2\%$ for the entire period, with a standard error of 3.6%. This reflects that time spent at recreational computing only increased 4.1 percent for older men while ESP increased 4.4 percent.

5.2 Impact on Labor Supply from Technology Change

The preceding subsection used shifts in time allocation to document that there has been rapid progress in technology associated with recreational computer use and video games. The question we now address is how this affects the willingness to work. From Section 3.4, equation (17) maps shifts in time allocations into shifts in leisure demand, holding constant the wage and marginal utility of consumption. The alternative mapping that assumes de-

[24] As a robustness exercise we instead assume no technological change in the weighted average of all other leisure activities. Our estimate of $(\eta_I - 1)\Delta \ln \theta_I$ is 43.5% with a standard error of 5.6%.

[25] We bootstrap our entire procedure to estimate the standard errors for our $(\eta_I - 1)\Delta \ln \theta_I$, using the ATUS replication weights.

clines in labor earnings generate equivalent declines in consumption is given by dividing this U_c constant prediction and by the denominator of equation (16). To quantify the wealth effect implicit in that denominator, note that $\epsilon(H/(1-H))$ is simply the Frisch elasticity of labor, which equals the Frisch elasticity of leisure times the ratio of leisure to non-leisure time.[26] We assume the Frisch elasticty of labor is equal to the inter-temporal elasticity of substitution in consumption, $\epsilon H/(1-H) = \rho$. The final term in the denominator of (16) is the ratio of labor income to consumption. We make a hand-to-mouth assumption and take this to be one. Therefore, the denominator is 2, and accounting for consumption changes reduces the U_c-constant effect on leisure by one half.

In addition to our estimates of the β's, $(\eta_I - 1)\Delta \ln \theta_I$, and time use data, we need two additional parameters to estimate how changes in leisure technology affect labor supply as given by (17). The first parameter is the Frisch elasticity of leisure, ϵ. The second is the average leisure-activity elasticity $\overline{\eta}$. The two parameters are related, as seen from equation (13). As a benchmark, we assume v enters linearly in U; that is, $\epsilon = \overline{\eta}$. In Section 6, we use price data and equation (18) to check the plausibility of the assumption $\epsilon = \overline{\eta}$, as well as to explore the robustness of our results to alternative choices of ϵ and $\overline{\eta}$.

We see from (17) that, if $\epsilon = \overline{\eta}$, the impact of technology on labor supply is independent of the level of ϵ. Although a higher ϵ implies a greater response to a given shift in technology, a higher $\overline{\eta}$ also implies that a given change in time allocation reflects a smaller increase in technology. When $\epsilon = \overline{\eta}$, the two effects cancel exactly, making our estimates independent of the level of ϵ.

Table 7 reports estimates for the shift in labor supply for our four demographic groups. To move from shifts in leisure demand (equation 17) to labor supply, we scale by the ratio of average leisure to average non-leisure time for each demographic group, with these averages based on the 2004-2015 ATUS. That is, $\Delta \ln n \approx -\Delta \ln H * (H/(1-H))$. Column 1 of Table 7 reports the estimated shift in labor supply assuming individuals are hand-to-mouth, while Column 2 reports the shift holding marginal consumption constant. As noted above, given the assumption that the labor Frisch equals the consumption Frisch, the Column 2 estimate is always twice Column 1. In both cases, wages are held constant, and hence the numbers should be interpreted as shifts in the labor supply *curve*.

To see how these estimates are constructed, consider younger men. Over the ATUS sample, the share of leisure devoted to computers (s_I) is 7.2 percent. As discussed in previous

[26] Technically, this is the elasticity of non-leisure time. In the data, non-leisure time is split between market work and home production (including child care, education, etc.). We assume that changes in leisure at the margin do not alter the share of non-leisure time devoted to market work. That is, additional leisure time is drawn from market work and home production proportionally. Thus a one percent decrease in non-leisure time is associated with a one percent decrease in both market work and home production.

Table 7: Impact of $\Delta\theta_I$ on Labor Supply

	Hand-to-Mouth	Full Insurance
Men 21-30	-1.53%	-3.06%
	(0.29%)	(0.58%)
Men 31-55	0.11%	0.22%
	(0.06%)	(0.13%)
Women 21-30	-0.40%	-0.80%
	(0.10%)	(0.21%)
Women 31-55	-0.16%	-0.32%
	(0.05%)	(0.09%)

Note: Table shows the shift in labor supply (wage constant) from $\Delta\theta_I$ for 2004-2007 to 2012-2015. Column 1 assumes one-to-one response of consumption to earnings (hand-to-mouth), Column 2 assumes no change in the marginal utility of consumption. Bootstrapped standard errors are in parentheses.

subsection, $\Delta \ln h_I - \frac{\beta_I}{\beta_{ESP}} \ln h_{ESP}$ is estimated at 38.7 percent. From equation (17) and setting $\epsilon = \bar{\eta}$, this implies a shift in leisure demand of 2.8 percent. Given that $H/(1-H)$ is 1.1, we have $\Delta \ln n = -3.1$ percent, which is the number reported in the table. If agent's are hand-to-mouth, this effect is reduced by half. Hence, our benchmark estimate is that the increase in computer leisure technology reduced labor supply for younger men by between 1.5 percent and 3.1 percent.

To put this shift in perspective, in the ATUS younger men exhibited an actual decline in market work between 2004 and 2015 of 6.7 percent.[27] Thus the shift in labor supply due to better computer technology constitutes 23 to 46 percent of the observed decline in hours for younger men in the ATUS since 2004. Keep in mind that our labor supply shifts holds the wage constant. How this shift translates into equilibrium wages versus market hours depends on the elasticity of labor demand. Given that younger men are a fairly small demographic group, and are likely highly substitutable with other workers, it is reasonable to assume that a relative shift in labor supply of younger men primarily affects their hours rather than wages. However, regardless of that mapping into equilibrium hours versus wages, the implied shift in labor supply is sizable even in the context of the large observed decline in market hours.

[27] For comparison the CPS data show a decline of 8.0 percent over those years.

A few other results are of note from Table 7. First, improved computer technology explains none of the decline in hours for older men. This stems from the facts that: (1) older men's share of time spent on computer activities is relatively small, and (2) they experienced little increase in the time spent on computer activities during the 2000s. These findings, coupled with the results for younger men in Row 1, suggest that increases in computer technology explain between 37 percent and 75 percent of the differential decline in hours worked for younger versus older men from 2004 to 2015.[28] Our estimates suggest that, absent the increase in computer technology, younger men would have exhibited a decline in market hours closer to that of older men.

Second, increased computer technology explains a decline in labor supply for younger women that is only one-fourth that for younger men, even though they also experienced a sizable percentage increase in computer leisure during the 2000s. This largely reflects the lower share of leisure younger women allocate to recreational computing; namely, 3.4 percent versus the 7.2 percent of younger men.

6 Robustness

Our base specification assumes that $\epsilon = \bar{\eta}$, which implies that leisure activities enter the utility aggregator U in an additive separable fashion. In this section we explore the plausibility of this assumption using price and expenditure data. We then examine the sensitivity of the results to alternative assumptions.

6.1 Estimating Technology Change from Prices and Expenditures

As discussed above, observed shifts in time allocation and the leisure Engel curves identify changes in technology up to the scaling parameter $\bar{\eta}$. Specifically, the leisure demand system allows us to measure $(\eta_I - 1)\Delta \ln \theta_I = (\bar{\eta}\beta_I - 1)\Delta \ln \theta_I$. To obtain a measure of $\bar{\eta}$, we need an independent measure of $\Delta \ln \theta_I$. We compute an estimate of $\Delta \ln \theta_I$ by using equation (12), assuming an interior solution, together with BLS price and expenditure data. The equation relates $\Delta \ln \theta_I$ to the difference in prices across technological vintages, $\Delta \ln p_I$, as well as the relative cost shares of goods (p_I) and time (wh_I) in the production of the leisure activity.

The relative prices of video games and equipment fell sharply during the 2000s. The BLS publishes a CPI for toys and games, which includes video games and equipment. The overall CPI increased 0.021 log points per year during the ATUS sample period of 2004-2015. Over

[28] From the 2004 to 2015 ATUS, younger and older men experienced respective declines in market hours of 6.7 versus 2.9 percent. CPS data from 2004 to 2015 also show a differential change of roughly 4 percent.

the same period, the annual rate for toys and games equaled -0.057 log points. For post-2008, the BLS has provided us the relative weight by year for the non-gaming component of "toys and games" as well as the price series for that non-gaming component. From this, we can infer that the price of the gaming component declined -0.127 log points per year. That is an annual price decline of 14.8 percent relative to the overall CPI. The CPI for computers and peripherals declined similarly, by 13.3 percent per year *relative* to the overall CPI. The BLS designs the CPI to be quality adjusted; that is, the price series ideally reflects the change in price holding quality constant. If the entry price of new models/vintages tracked the overall CPI, then the annual relative decline in the category's CPI captures the relative price across introductions of newer vintages.[29] The log price difference across annual vintages then should reflect the rate of increase in the overall CPI relative to a CPI for computers, peripherals, and video games. We put this rate, perhaps conservatively, at 13.3 percent per year.

We showed in (12) that one can recover $\Delta \ln \theta_I$ based on the relative price change for computer leisure goods together with these goods cost share in the activity. We take the marginal purchaser to be the average person in our sample. We deflate nominal quantities by the PCE deflator in 2009 dollars. Using the Consumer Expenditure Survey (CE), we break out expenditure on computers, video games, and peripherals. Reported expenditure on these goods in the CE averaged \$464 for 2004 to 2014 (in 2009 dollars), where we average over households with a member between the ages 21 and 55. Time spent on recreational computing for this period averaged 124 hours per year, where again we average over all respondents ages 21 to 55. From the CPS, the median real wage for the period for employed individuals ages 21 to 55 is \$17.9. Assuming a marginal tax rate of 25 percent, the after-tax wage is \$13.4. Using this as the opportunity cost of time, the time input into computers and gaming is \$1,660. Hence, an estimate of the goods-to-time cost ratio is 0.28. From equation (12), and a price decline of 13.5 percent per year, this implies annual technological progress for computers and video games of 3.7 percent a year.

As context for the 3.7 annual growth in computer and gaming technology, nominal expenditure on computers and peripherals by households with younger men increased at an annual rate of 8.6 percent (CE data). Deflating by the CPI price index for computers and peripherals, this represents a real increase of 20.2 percent *per annum*.[30] While all of the

[29] Tracking prices across vintages is complicated by the alternate varieties and features that are introduced with new models. For reference, the original Xbox was introduced in 2001 retailing for \$299.99. The next generation Xbox 360 arrived in 2005, with the "core" system selling for \$299.99 and the "bundle" for \$399.99. The Xbox One entered in 2013 at \$499.99, which included a Kinect sensor that sold separately for \$150.

[30] For the sample period 2012-2014, average nominal expenditure is \$571. The corresponding figure for 2004-2006 is \$288, representing an annual nominal growth rate of 8.6 percent. The decline in the CPI Price Index for computers and peripherals, also calculated as the difference in three-year averages, is 11.6 percent. Thus real expenditures increased at an annual rate of 20.2 percent.

expenditure on computers and peripherals is not solely for leisure, it does provide a sense of the substantial increase in computer and gaming hardware in the typical household. This naturally should increase the return on the time spent computing and gaming, which is reflected in our estimated $\Delta \ln \theta_I$.

Comparing our $(\eta_I - 1)\Delta \ln \theta_I = 4.8$ percent per year number, obtained from the shifts in time allocation, to the $\Delta \ln \theta_I = 3.7$ percent per year from price data, yields an η_I of 2.30. Using our estimated Engel curve $\hat{\beta}_I = 2.46$ and $\beta_I = \eta_I/\bar{\eta}$, we obtain $\bar{\eta} \approx 0.93$. Given this estimate, our benchmark assumption that $\epsilon = \bar{\eta}$ implies a leisure Frisch elasticity of 0.93. This calculation provides a sense of the magnitude of $\Delta \ln \theta_I$ from price and expenditure data, and hence the scale parameter $\bar{\eta}$. Given the assumptions and data challenges involved, it should be viewed as a rough guide rather than a firm estimate. For this reason, in the next subsection we explore how our results vary with alternative values of $\bar{\eta}$ and ϵ.

6.2 Sensitivity of Results to ϵ and $\bar{\eta}$

In Section 5.2, we assumed that $\epsilon = \bar{\eta}$. In doing so, we did not need to specify a specific value for either variable. However, the size of the induced shift in labor supply more generally depends on the values of ϵ and $\bar{\eta}$. Equation (17) indicates exactly how our benchmark result varies with alternative values of these two parameters, showing that the magnitude is scaled by the factor $\left[\frac{\epsilon\beta_I - 1}{\bar{\eta}\beta_I - 1}\right]$. Here we explore robustness of the implied impact on labor supply to varying both ϵ and $\bar{\eta}$.

There is an extensive literature estimating the Frisch labor supply elasticity. Recall that as leisure is roughly half the discretionary time in our framework, the leisure Frisch is approximately equal to the labor Frisch. Moreover, the relevant elasticity for our framework is the combination of the extensive and intensive margins. Hall (2009) surveys the literature estimating the intensive margin Frisch. He takes its value to be in the range of 0.7, with that choice especially influenced by Pistaferri (2003)'s estimate of 0.71. Chetty et al. (2013) similarly survey a number of estimates of the intensive margin Frisch and arrive at a somewhat smaller consensus value of 0.54. Chetty et al. (2013) also survey several quasi-experimental estimates of the extensive-margin Frisch elasticity. They put the extensive elasticity, at 0.32. Several authors have produced structural estimates of the Frisch elasticity at the extensive margin. These suggest modestly higher elasticities, in the range of 0.4 to 0.7. (See Gourio and Noual (2009), Mustre-del-Río (2015), and Park (2017).) Based on this literature, we treat the *combined* Frisch, reflecting both the intensive and extensive responses, to be in the neighborhood of one. So, to examine robustness, we let its value vary across $\{0.75, 1.0, 1.25\}$.

From the calculations in the last subsection, we arrived at 0.9 as a plausible value for

Table 8: Robustness of ϵ and $\bar{\eta}$ on Labor Supply Impacts of Young Men

	$\bar{\eta} = 0.75$	$\bar{\eta} = 1.0$	$\bar{\eta} = 1.25$
$\epsilon = 0.75$	-3.1%	-1.8%	-1.2%
$\epsilon = 1.00$	-5.3%	-3.1%	-2.2%
$\epsilon = 1.25$	-7.5%	-4.4%	-3.1%

Note: Table shows the shift in labor supply (wage constant) from $(\eta_I - 1)\Delta\theta_I$ for 2004-2007 to 2012-2015 for younger men. The entries in this table corresponds to row 1 of Table 7. The table shows the robustness of those results to various alternate values of ϵ and $\bar{\eta}$.

$\bar{\eta}$, though with admittedly some uncertainty attached to that calculation. We consider the same range of values for $\bar{\eta}$ as taken for ϵ, that is, $\{0.75, 1.0, 1.25\}$.

The implied change in labor supply of younger men due to changes in leisure technology is reported in Table 8 for these alternative values for parameters ϵ and $\bar{\eta}$. For ease of exposition, we only show the results holding the marginal utility of consumption constant. As above, the hand-to-mouth estimates are approximately one-half the constant marginal utility of consumption estimates. Recall that our benchmark sets $\epsilon = \bar{\eta}$. Hence, the diagonal of the table replicates our baseline estimate of a 3 percent decline in labor supply.

Fixing ϵ, we see that an increase in $\bar{\eta}$ reduces the implied shift in labor supply. For example, holding ϵ constant at 1.0, the shift in labor supply ranges from -5.3% to -2.2% as $\bar{\eta}$ increases from 0.75 to 1.25. Recall from equation (13) that the Frisch elasticity can be decomposed into $\bar{\eta}$, the average elasticity within v, and the additional curvature due to the leisure aggregator U. As we hold ϵ constant and increase $\bar{\eta}$, we increase the curvature of U, which lowers the responsiveness of leisure to an increase in technology.

Reading down a column, fixing $\bar{\eta}$, a higher Frisch elasticity increases the implied shift in labor supply. For example, fixing $\bar{\eta} = 1.0$, the implied shift in labor ranges from -1.8% to -4.4% as the Frisch elasticity varies between 0.75 and 1.25. While it is clear that the relative magnitude of ϵ to $\bar{\eta}$ plays an important role in the quantitative impact of computer and gaming technology on labor supply of younger men, for a wide range of these parameters the estimated impact remains quite substantial.

Table 9: Share of Younger Men and Women Living With Parent or Close Relative

	Men 21-30	Women 21-30
2000	0.23	0.20
2007	0.27	0.26
2010	0.31	0.29
2015	0.35	0.34
Change 2000-15	0.12	0.14

Note: Table shows the fraction of men and women ages 21-30 cohabitating with their parents/step-parents or other close relatives (siblings, grandparents, etc.). Data are from the American Community Survey.

7 Younger Men's Consumption and Well Being

We find that the impact of innovations to recreational computing on younger mens' labor supply depends on how well their consumption is insulated - if they sacrifice earnings for gaming. In this section, we show that younger individuals - particularly younger men - receive substantial inter-family transfers when they do not work.

7.1 Trends in Cohabitation and Consumption

Table 9 documents cohabitation patterns of younger men and women as seen from the 2000 Census and the 2001-2015 American Community Surveys (ACS).[31] The first column shows the trend in younger men living in a household where a parent, step-parent, or other close relative (sibling, grandparent, uncle, aunt) is the household head. In 2000, 23 percent of younger men lived with a close relative. By 2015 that fraction was 35 percent, with the change driven mostly by an increase in living with parents. From column 2, younger women are less likely to live with parents, but experienced a similar upward trend during the 2000s.

Table 10 shows cohabitation patterns for younger men by employment status, pooling data for 2000-2003 and 2012-2015. We summarize the key takeaways. (1) Non-employed younger men are 20 percentage points more likely, to live with parents/relatives. In 2012-2015, 67 percent of those not working lived with a parent or close relative, with only 12 percent living on their own. (2) Between 2000-2003 and 2012-2015, periods differing by

[31]The 2000 Census and subsequent ACS contain comparable questions on a respondent's relationship to the household head. A head is the person (or persons) that owns or rents the housing. As with the ATUS sample, we exclude full-time students ages 25 or less. We also exclude those residing in group quarters. See the Online Appendix for added detail on the Census/ACS samples.

Table 10: Changes in Younger Men's Household Status

	2000-2003 Data		2012-2015 Data	
Living Status	Employed	Non-Emp.	Employed	Non-Emp.
Head: Single	0.25	0.19	0.23	0.12
Head: Live with Spouse/Partner	0.41	0.26	0.28	0.12
Not Head: Live with Parent/Close Rel.	0.26	0.46	0.37	0.67
Not Head: Live with Others	0.09	0.08	0.12	0.09

Table shows the fraction of younger men in each cohabiting arrangement by employment status. Data are from the American Community Survey (ACS). Full-time students ages less than 25 are excluded. We classify as household heads anyone who reports being the household head, the spouse of the household head, or the unmarried domestic partner of the household head. (Household head in ACS refers to the individual that owns or rents the housing unit.)

only 12 years on average, there was a dramatic increase of 13 percentage points (nearly 50 percent) in living with parents and other close relatives. In the early 2000s, 26 percent of employed and 46 percent of non-employed younger men lived with a parent or close relative. By 2012-2015, those shares were 37 percent and 67 percent. The importance of cohabiting with parents has been emphasized in the business-cycle context by Kaplan (2012) and Dyrda et al. (2012). We document that it is also relevant for the longer-run decline in employment of younger men.

By 2012-2015, only 12 percent of non-working younger men are married or live with a partner. A similarly small fraction report living in a household with a child. Given these younger men are neither married nor have children in the household, government programs are not a major source of income. Younger single men without children do not receive welfare programs like SNAP. Their lack of work experience means many do not receive unemployment benefits. Disability take-up is also rare for this age group. Thus parents and other relatives are the more likely source for support, especially housing.

Younger men living on their own may still receive support from their parents. To examine this, we use biannual surveys from the Panel Study of Income Dynamics (PSID) for 2001 to 2013. From the PSID it is possible to see transfers in the form of help from relatives. (A fuller description of our PSID sample can be found in the Online Appendix.) We highlight a few takeaways. First, help from relatives is still fairly common for younger men that do not live with relatives, with about 20 percent of these households reporting such help. But these transfers are typically small, averaging (including zeros) only 1.5 percent of those households' average earnings. Second, government transfers are fairly small for households headed by younger men. Government transfers (e.g., unemployment benefits, SSI benefits) averaged 2.9

percent of household earnings for these households, while tax credits (EITC, child credits, etc.) averaged another 1.9 percent. Finally, government transfers are much more important for households where younger men live with parents or other relatives. Across the seven PSID waves, government transfers and credits represented 15.7 percent of average earnings for these households. They also increased substantially over time, equaling 22.1 percent of average earnings for such households by the 2013 survey. Social security benefits are easily the most important government payment for these households, both in terms of level and trend growth. These government payments presumably contribute toward spending by younger men in these households, even if they are not the direct beneficiary.

In the Online Appendix, we use PSID measures to track expenditures in households with younger men, versus those with older men, in light of their differential trends in hours worked documented above. The analysis is imperfect, in that expenditures are measured at the household level while our analysis on employment and hours concerns individuals. We take the standard approach of deflating household expenditures by a measure of household scale (equivalence units), cognizant that this imposes the assumption that expenditures are split equally between the parent and the dependent. The PSID data indicate that younger men's consumption, adjusted for household size, does not decline relative to households containing older men. In particular, households containing a younger man experienced a decline in after-tax income of 6.6 percent between 2000 and 2012, but recorded less than a one percent decline in consumption. Households containing men age 31-55 experienced a smaller decline in income but a larger decline in expenditure. We view the consumption data as reinforcing the cohabitation trends as evidence that parents and close relatives are providing significant consumption insurance to younger men during the 2000s.

7.2 Trends in Well-Being

Before concluding, we turn to data from the General Social Survey (GSS) to examine trends in reported life satisfaction for younger men relative to other groups. The GSS assesses attitudes and beliefs of US residents. The GSS has consistently asked individuals the following question: "Taken together, how would you say things are going these days – would you say that you are very happy, pretty happy, or not too happy?" We create a happiness index that equals 1 if an individual reports being either "very happy" or "pretty happy," and equals 0 otherwise. As with the ATUS, we pool waves of the GSS index, given the survey's modest sample size.[32] We examine three time periods: 2001 to 2005, 2006 to 2010, and 2011 to 2015.

Table 11 tracks happiness measures for younger versus older men, first for all education

[32]The survey is biannual and nationally representative. Each GSS wave has 2,000 to 4,000 respondents.

Table 11: Reported Happiness

	Fraction Reporting "Very Happy" or "Pretty Happy"				
	(1) Pooled 2001-2005	(2) Pooled 2006-2010	(3) Pooled 2011-2015	Diff (3)-(1)	p-value of difference
Men, Ed = All, 21-30	0.839 (n=249)	0.854 (n=507)	0.892 (n=343)	0.053	0.060
Men, Ed = All, 31-55	0.886 (n=630)	0.854 (n=1,528)	0.847 (n=903)	-0.039	0.031
Men, Ed < 16, 21-30	0.813 (n=193)	0.828 (n=372)	0.881 (n=244)	0.068	0.048
Men, Ed < 16, 31-55	0.883 (n=426)	0.828 (n=1,043)	0.813 (n=594)	-0.069	0.023

Note: Data from General Social Survey. See text for details.

groups, then excluding those with 4 or more years of college. The happiness of younger men actually increased by 5 percentage points since the early 2000s, from 84 to 89 percent, despite their sharp decline in employment. This stands in contrast to the pattern for older workers, for whom measured happiness fell by 4 percentage points. In the early 2000s, older men reported being happier than did their younger counterparts. That relationship flipped by 2011-2015. The contrast is even more striking restricting attention to those with less than a 4-year degree (rows 3 and 4). Younger men without a 4-year degree show a 7 percentage point increase in happiness, compared to a 7 percentage point decline among their older counterparts. Measures of well being for older workers has been studied recently by Case and Deaton (2015). Table 11 adds to this literature showing, by contrast, that younger men experienced a rise, rather than decline, in measured happiness over the past 15 years.

While by no means conclusive, these results are consistent with computer technology broadly, and video games in particular, increasing the value of leisure for younger workers.[33]

[33] Krueger (2017) uses ATUS self-reported well-being measures to compare the emotional experience of young men across various leisure activities. He finds that younger individuals report greater happiness when playing video games relative to watching TV. He also finds that video game playing is a social activity for younger men—70% of time playing involves interacting with others in person or virtually.

8 Conclusion

In this paper we develop a leisure demand system that parallels that typically considered for consumption expenditures. This allows us to estimate how leisure activities vary with one's total leisure time, generating activity-specific leisure Engel curves. Our framework also provides a means for assessing how much improvements in leisure technologies can affect individual's labor supply. We show that such innovations are likely to reduce labor supply much more if they affect leisure luxuries. Estimating our leisure demand system based on leisure differences across time, states, industries, and education groups during the 2000s, we find that recreational computer activities in general, and video gaming especially, are strong leisure luxuries for younger men. We estimate that younger men respond to a 1 percent increase in total leisure by increasing recreational computer time by 2.5 percent. For other groups – younger women, older men, and older women – recreational computing is only modestly a leisure luxury.

Using our estimated leisure demand system, together with detailed time use data from the American Community Survey, we can identify the relative increase in computer and video game technology during the 2000s. As of 2015, men between the ages of 21 and 30 allocated 5.2 hours per week to recreational computer activities, 3.4 hours going specifically to video gaming. For younger men recreational computer time increased by 45 percent during the 2004-2015 period, while total leisure time increased by only 4 percent. Our estimated leisure demand system predicts that recreational computer time would have increased by 8 percent if younger men had remained on their original leisure Engel curve. We can attribute the much greater increase in younger men's computer time to a sizable improvement in technology for computer and video gaming, an improvement we would expect given CPI-measured declines in relative prices for computer and video games.

We estimate that technology growth for recreational computer activities, by increasing the marginal value of leisure, accounts for 23 to 46 percent of the decline in market work for younger men during the 2000s. Our estimates also suggest that technology growth can explain as much as three-quarters of their 4 percent greater decline in hours relative to men ages 31-55. We estimate that improved computer and gaming technology explains a small decline in market work for younger women, but had no impact for older men and women.

Presumably innovations to gaming and computer leisure permeate national borders. So, if these innovations affected younger men's labor supply in the U.S., then we should expect an impact in other countries. Figure 5 plots the trends in employment to population rates for men ages 25 to 29, relative to that for men ages 30 to 54, for both the U.S. and the

Figure 5: Employment Rates for Men 25-29 versus Men 30-54 since 2000, U.S. and OECD

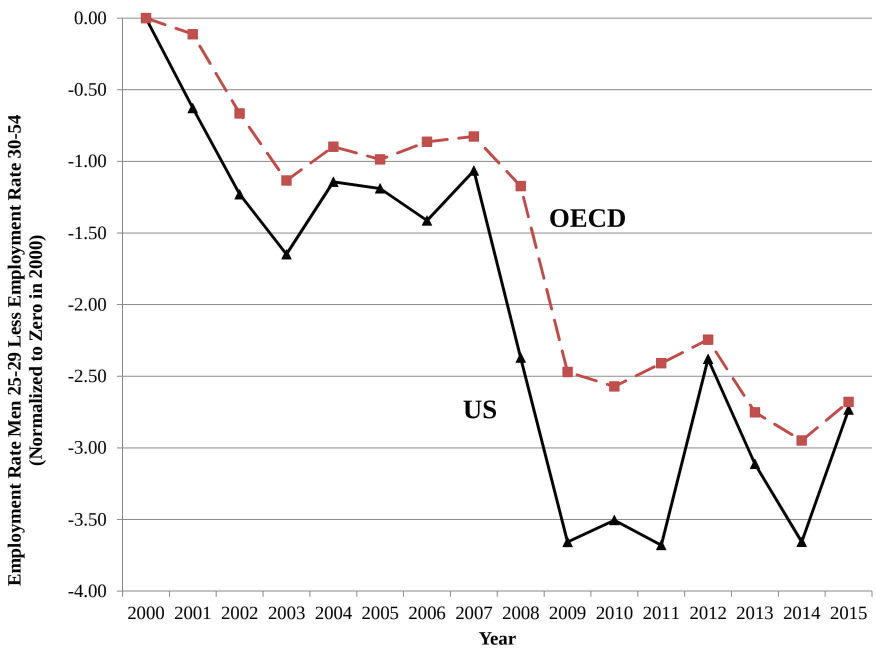

Note: Figure shows the employment to population rate for men 25-29 minus that rate for men 30-54 for the U.S. and for OECD countries. Data are from OECD.Stat. The differential is expressed in percentage points. It is normalized to zero in 2000, so the series are relative to 2000.

OECD.[34] We see that OECD countries displayed the same decline in relative employment for younger men, -2.7 percentage points, as the U.S. decline. Compared to the U.S., relative employment for younger men fell less sharply during the Great Recession. But, while it has partially rebounded in the U.S., it has continued a slight decline in the OECD. Some OECD countries (notably the PIIGS) experienced particularly depressed labor markets over this period, disproportionately affecting younger workers. But, if we restrict attention to Canada, the U.K., and Australia, countries arguably more comparable to the U.S., we still see declines in relative employment for younger men since 2000 that mirror the U.S. experience: A relative decline of nearly 3 percentage points for Canada; and about 4 and 2 points, respectively, for the U.K. and Australia.

Our focus on computer leisure was driven by the sharp shifts we see in younger men's time use since the early 2000's and by a prior, confirmed in the data, that computer time

[34]Data are from OECD.Stat; it dictates the age breaks. Many countries report time use data, but not at sufficient detail to isolate changes in gaming and computer leisure time over time.

is a leisure luxury for younger men. However, the methodology developed here is applicable to innovations to any leisure activity. In particular, the predicted impact on labor supply will be larger if that activity is a leisure luxury. Aguiar and Hurst (2007) show from U.S. time-use data that time spent watching television increased by 7 hours per week on average for men and women ages 21 to 65 from the mid 1960's to the early 2000's. That constituted more than 100 percent of a sizable increase in total leisure over that period of 5.5 hours per week. We believe it is plausible that the increase in total leisure over those 40 plus years partially reflected the development of television programming, especially given that we estimate watching television is a clear leisure luxury for older men.

In this paper, we have developed a methodology for measuring changes in the return to leisure which is a component of an individual's reservation wage. The extent that increases in leisure technology align empirically with reductions in market hours depends, of course, on how those leisure technology shifts happen to coincide with factors shifting labor demand. In periods where labor demand and reservations wages are both increasing (like during the 1970s and 1980s when the quality of television expanded rapidly), increases in leisure technology may not correspond with declines in employment. However, during the 2000s, market wage growth was declining, reflecting declining labor demand, while the reservation wage was arguably increasing. For any individual, an increase in leisure technology is more likely to result in declining employment when the market wage is close to the reservation wage.

Finally, our framework is static. However, innovations to computer and gaming leisure may have dynamic effects on labor supply. It is possible that individuals develop a habit (or addiction) for such activities. Certainly individuals build "leisure capital" in the form of physical equipment, but especially human skills, that enhances enjoyment from gaming. Thus negative shocks to labor demand could have a persistent negative impact on labor supply via individuals first increasing their computer leisure, then developing a taste or skills for the activity. Such dynamic consideration may be a source of hysteresis in labor market conditions resulting from downturns, such as the Great Recession. We leave these considerations to future work.

References

Aguiar, Mark and Erik Hurst, "Measuring Trends in Leisure: The Allocation of Time over Five Decades," *Quarterly Journal of Economics*, 2007, *122*, 969–1006.

Becker, Gary, "A Theory of the Allocation of Time," *Quarterly Journal of Economics*, 1965, *75*, 493–517.

Browning, Martin and Thomas Crossley, "Luxuries are Easier to Postpone: A Proof," *Journal of Political Economy*, 2000, *108*, 1022–6.

_ , Angus Deaton, and Margaret Irish, "A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle," *Econometrica*, 1985, *53*, 503–543.

Case, Anne and Angus Deaton, "Suicide, Age, and Wellbeing: An Empirical Investigation," *NBER Working Paper*, 2015, *21279.*

Chang, Yongsung and Sung-Bin Kim, "From Individual to Aggregate Labor Supply: A Quantitative Analysis based on a Heterogeneous-Agent Macroeconomy," *International Economic Review*, 2006, *47*, 1–27.

Charles, Kerwin, Erik Hurst, and Matthew Notowidigdo, "Housing Booms, Manufacturing Decline, and Labor Market Outcomes," *NBER Working Paper*, 2016, *18949.*

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber, "Does Indivisible Labor Explain the Difference Between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities," *NBER Macroeconomics Annual 2012*, 2013, *27.*

Crossley, Thomas and Hamish Low, "Is the Elasticity of Intertemporal Substitution Constant?," *Journal of the European Economic Association*, 2011, *9*, 87–105.

Deaton, Angus and John Muellbauer, "An Almost Ideal Demand System," *American Economic Review*, 1980, *70*, 312–326.

Dyrda, Sebastian, Greg Kaplan, and José-Víctor Ríos-Rull, "Business Cycles and Household Formation: The Micor vs the Macro Labor Elasticity," *working paper*, 2012.

Gourio, Francios and Pierre-Alexandre Noual, "The Marginal Worker and the Aggregate Elasticity of Labor Supply," *working paper*, 2009.

Greenwood, Jeremy and Guillaume Vandenbroucke, "Hours Worked (Long-Run Trends)," in L.E. Blume and S.N. Durlauf, eds., *The New Palgrave Dictionary of Economics*, Palgrave, 2008, chapter Hours Worked (Long-Run Trends).

Hall, Robert, "Reconciling Cyclical Movements in the Marginal Value of Time and the Marginal Product of Labor," *Journal of Political Economy*, 2009, *117*, 281–323.

Kaplan, Greg, "Moving Back Home: Insurance Against Labor Market Risk," *Journal of Political Economy*, 2012, *120*, 446–512.

Kopecky, Karen, "The Trend in Retirement," *International Economic Review*, 2011, *52*, 287–316.

Krueger, Alan, "Where Have All the Workers Gone," *Brookings Papers on Economic Activity*, 2017, *Fall 2017.*

Mian, Atif and Amir Sufi, "What Explains the 2007-2009 Drop in Employment?," *Econometrica*, 2014, *82*, 2197–2223.

Mincer, Jacob, "Labor Force Participation of Married Women: A Study of Labor Supply," in "Aspects of Labor Economics," Princeton University Press, 1962.

Mustre-del-Río, José, "Wealth and Labor Supply Heterogeneity," *Review of Economic Dynamics*, 2015, *18*, 619–634.

Park, Choonsung, "Consumption, Reservation Wages, and Aggregate Labor Supply," *working paper*, 2017.

Pistaferri, Luigi, "Anticipated and Unanticipated Wage Changes, Wage Risk, and Intertemporal Labor Supply," *Journal of Labor Economics*, 2003, *21*, 729–54.

Vandenbroucke, Guillaume, "Trends in hours: The U.S. from 1900-1950," *Journal of Economic Dynamics & Control*, 2009, *33*, 237–249.

Online Appendix for “Leisure Luxuries and the Labor Supply of Young Men”

A1 Additional Tables and Figures

Table A1: Leisure Engel Curves of Younger Men: $\hat{\gamma}_i$

	(1)	(2)	(3)	(4)
Recreational Computer	0.11	0.11	0.09	0.05
	(0.01)	(0.01)	(0.01)	(0.02)
Video Games	0.09	0.08	0.07	0.05
	(0.01)	(0.01)	(0.01)	(0.01)
TV/Movies/Netflix	0.05	0.02	0.04	0.04
	(0.02)	(0.02)	(0.02)	(0.02)
Socializing	-0.04	-0.03	-0.06	0.00
	(0.01)	(0.01)	(0.01)	(0.02)
ESP	-0.12	-0.11	-0.09	-0.09
	(0.02)	(0.02)	(0.02)	(0.03)
Other Leisure	0.00	0.02	0.02	0.01
	(0.01)	(0.01)	(0.01)	(0.01)
Fixed Effects:				
Time Period	✓	✓	✓	✓
Education		✓	✓	✓
Geographic			✓	✓
Industry				✓
Number of Cells	242	242	242	242
Number of Individuals	6,250	6,250	6,250	6,250

Note: Estimated $\hat{\gamma}_i$ from AIDS specification (19). These estimates are used to construct $\hat{\beta}_i$ reported in Table 5. An observation is a time-gender-age-education-industry-state group cell. Bootstrapped standard errors are in parentheses.

Figure A1: Distribution of Leisure Time for Young Men

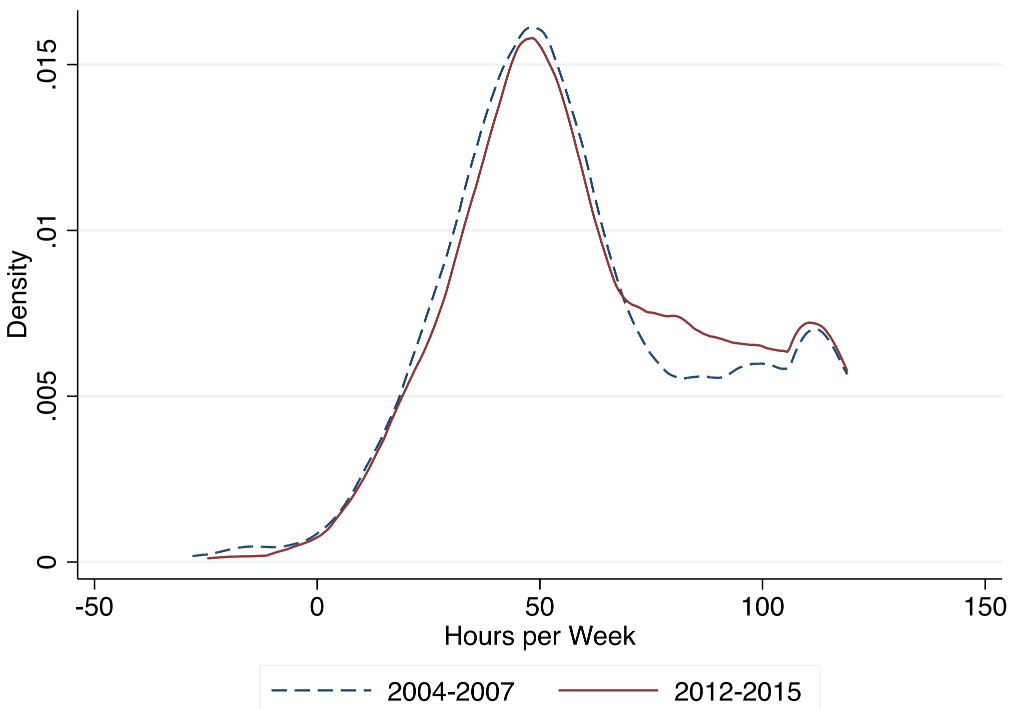

Note: Figure shows kernel density of leisure time for younger men.

Figure A2: Distribution of Recreational Computing Time for Young Men

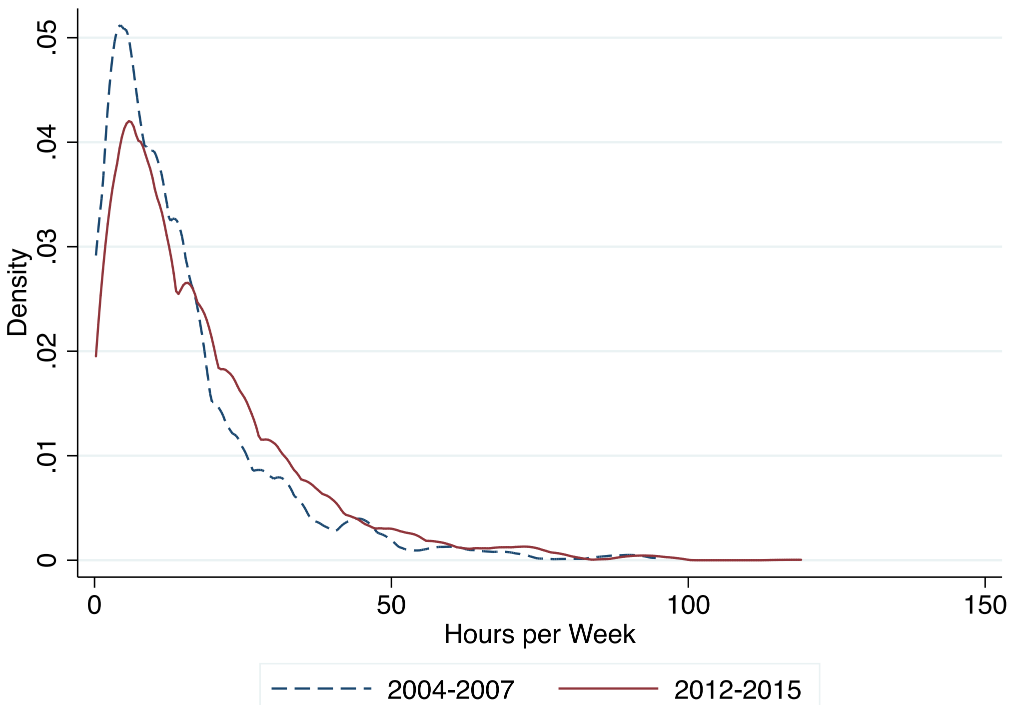

Note: Figure shows kernel density of recreational computing time for younger men conditional on strictly positive time. The share of younger men with zero computing time is 0.23 for the 2004-07 ATUS sample and 0.28 for the 2012-15 ATUS sample.

A2 Data Appendix

We primarily use three data sets in the analysis. In this section, we provide information – including sample restrictions – for the American Time Use Survey, the Current Population Survey, and the Census/American Community Survey. For some supplementary analysis, we also use data from the Panel Study of Income Dynamics, the General Social Survey and the BLS Price data. The main text discusses our use of these datasets.

A2.1 American Time Use Survey (ATUS)

The bulk of our analysis is based on the 2004 to 2015 waves of the American Time Use Survey (ATUS). The ATUS is conducted by the U.S. Bureau of Labor Statistics (BLS), with individuals drawn from the exiting sample of the CPS. (We download the ATUS data directly from the BLS website.) Individuals are sampled approximately 3 months after completion of their final CPS survey. At the time of the ATUS, the BLS updates the respondent's employment and demographic information. The time-use data reflect a 24-hour diary where respondents report activities from the previous day broken by 15 minute intervals. Survey personnel then classify each activity to a specific one of over 400 detailed categories. We omit a few minor time categories, such as own health and a catch-all "uncategorized" activity.

The time diaries are designed to measure an individual's primary task. It measures secondary tasks less well. For example, consider someone who commutes for a half hour on the subway, reading a book during their commute. The survey will prompt the individual to only report the primary activity, which would likely be commuting. However, if the individual lists multiple activities as their primary activity, those activities get allocated an equal portion of that time interval. Continuing the example, if someone reported both commuting and reading were primary activities, 15 minutes would get allocated to commuting and 15 minutes to reading. This preserves that each individual's has total reported time of 24 hours. So it is likely that the less primary of multi-tasking activities are underreported. This may be relevant to some types of recreational computer activities, like engaging in social media.

Time spent at market work in the ATUS diaries differs from that reported in the CPS March supplements. The time diary includes commuting time. It is also the amount worked for one 24-hour period, rather than a recall estimate of hours worked in a "usual" week.

We restrict the sample to those ages 21 to 55. We exclude individuals in the military and full-time students ages 24 and under. Status as a full-time student is only consistently asked in the March Supplement for those ages 24 and under. The last four columns of Table A2 report by year the number of individuals in our ATUS sample.

A2.2 Current Population Survey (CPS)

We downloaded the 1977-2016 March Annual Social and Economic Supplements to the CPS directly from the IPUMS CPS website (https://cps.ipums.org/cps/index.shtml). We restrict the sample to ages 21 to 55, again excluding individuals in the military and full-time students ages 24 and under. (Status as a full-time student is only consistently asked in the March Supplement for those 24 and under.)

Table A2: Annual Sample Sizes from CPS and ATUS

	CPS				ATUS			
	All 21-55	Men 21-30	Men 31-55	Women 21-30	All 21-55	Men 21-30	Men 31-55	Women 21-30
2000	63,996	7,611	23,394	8,101				
2001	104,078	11,246	38,531	12,684				
2002	103,344	10,742	38,400	12,388				
2003	102,548	10,724	38,146	12,027				
2004	100,460	10,486	37,178	11,913	8,579	656	3,107	963
2005	98,663	10,571	36,273	11,815	8,255	630	2,947	957
2006	97,445	10,716	35,638	11,783	7,982	606	2,845	881
2007	96,147	10,644	35,212	11,633	7,513	615	2,736	816
2008	95,437	10,641	34,782	11,599	7,734	642	2,836	842
2009	95,976	10,660	35,074	11,730	7,941	617	2,901	920
2010	96,577	10,908	35,266	11,929	7,987	636	2,940	923
2011	93,703	10,585	34,100	11,752	7,367	546	2,719	889
2012	91,397	10,193	33,345	11,327	7,216	589	2,691	781
2013	91,356	10,414	33,379	11,211	6,457	517	2,395	736
2014	89,142	10,217	32,383	11,011	6,494	540	2,389	729
2015	88,311	10,236	31,901	11,171	6,073	474	2,243	673
2016	81,905	9,447	29,502	10,195				

Note: Table shows sample sizes for our analysis samples from the CPS (first four columns) and the ATUS (last four columns) by year. Our ATUS sample only includes years between 2004 and 2015. See text for exact sample restrictions.

Our CPS series focus on hours and employment. We define those who are employed as anyone who reports working last week (empstat = 10) and anyone who has a job but did not work last week (empstat = 12). Employment status is measured as of the survey. For example, respondents in the 2016 March Supplement report information about whether they were working in March of 2016. Hours worked are reported retrospectively. Survey respondents in year t report (1) how many weeks worked during the prior calendar year and (2) the hours per week they usually worked during the prior year. We construct annual hours worked by multiplying weeks worked during the prior year by the usual weekly hours worked during the prior year. We also document the extent to which individuals did not work during the prior year. We define not working during the prior year as survey respondents who report working zero weeks during the prior year.

The first four columns of Table A2 show the sample sizes for our full sample, younger men sample, younger women sample, and older men sample for 2000 to 2016.

A2.3 Census/American Community Survey (ACS)

We use data from the 2000 Census and the 2001-2015 American Community Surveys (ACS) to validate the patterns in market hours from the ATUS and CPS samples and to test robustness of our sample restrictions. Finally, and most importantly, we use the Census/ACS data to measure the trends in cohabitation shown in Section 7.1 of the paper.

The ACS's are annual surveys starting in 2001 detailing socio-economic information for a large sample of Americans. It is conducted by the U.S. Census Bureau and asks questions similar to the traditional Census long form. As a result, the 2000 Census and 2001-2015 ACS's ask a comparable set of questions and have a similar sampling frame. We download the data directly from the IPUMS website (https://usa.ipums.org/usa/). We include individuals between the ages of 21 and 55. To make the data consistent with our CPS sample, we exclude those in the military, those living in group quarters, and full-time students ages 24 and under. In a robustness specification, discussed below, we further exclude all full-time students regardless of age. The Census/ACS samples are large: For 2000, the sample is just under 6.5 million individuals; for 2001 to 2004, it ranges between 500,000 and 560,000 individuals per year; and, starting in 2005, it is about 1.25 million individuals per year. Its large sample is a key advantage of the Census/ACS data.

As with the CPS, we measure annual hours worked by multiplying weeks worked last year times usual hours worked. Here there is one key difference between the CPS and Census/ACS data. In the Census/ACS individuals report weeks worked over the last 12 months, not over the last calendar year. Given that the Census/ACS conducts interviews throughout the year, and because researchers do not have access to the month a respondent was surveyed, there is no direct mapping from this Census/ACS annual measure to the calendar year. Absent a solution, we map survey responses within year t to annual hours worked in year t. For example, 2015 respondents yield hours worked for 2015. As a result, some caution is needed when comparing trends in annual hours worked between the CPS and Census/ACS data.

A3 Trends in Employment and Hours from the CPS

In this section we compare shifts in employment and hours for younger men to other demographic groups during the 2000s using data from the March Current Population Survey (CPS). Table A3 Panel (a) reports annual hours worked for men and women at four points over the last 15 years. Panel (b) reports the same for those with less than a college degree. From 2000 to 2015, annual hours worked by younger men declined by 203 hours (11.8 percent) while the decline for older men was 163 hours (8.2 percent). If we restrict the samples to those without a 4-year degree, the respective declines are 14.4 versus 10.2 percent. We see from Table A3 that market work hours also declined for both younger and older women during the 2000s. But these declines were one-third to one-half that of their male counterparts. In sum, younger men, especially those without a 4-year degree, exhibited the largest decline in hours worked during the 2000s.

Table A3: Annual Market Hours Worked

(a) All Education

	Men		Women	
Year	21-30	31-55	21-30	31-55
2000	1,829	2,050	1,407	1,452
2007	1,728	1,964	1,355	1,429
2010	1,519	1,796	1,218	1,351
2015	1,626	1,887	1,312	1,398
Change 2000-15	-203	-163	-95	-54
Log Change 2000-15 ($\times 100$)	-11.8	-8.2	-7.0	-3.8

(b) Education < 16

	Men		Women	
Year	21-30	31-55	21-30	31-55
2000	1,801	1,953	1,311	1,397
2007	1,691	1,859	1,227	1,346
2010	1,436	1,658	1,080	1,241
2015	1,559	1,763	1,167	1,258
Change 2000-15	-242	-190	-144	-139
Log Change 2000-15 ($\times 100$)	-14.4	-10.2	-11.7	-10.5

Note: Data are from the March CPS. Annual hours equal last year's weeks worked multiplied by usual weekly hours. Year t hours refer to hours worked by year $t+1$ respondents. Full-time students less than age 25 are excluded.

A4 Robustness of Trends in Market Work, Census/ACS

We use Census/ACS data to explore trends in hours across demographic groups that parallel those reported above. The Census and ACS data capture full-time school enrollment over the 2000-2015 period for all individuals, not just those under age 25; so the Census/ACS data allow us to explore the robustness to excluding all full-time students, not just those under age 25. Panel A of Table A4 is analogous to Table A3, but based on the Census/ACS imposing the same sample restrictions as done with the CPS. In particular, panel (a) excludes only those full-time students under age 25. The CPS and Census/ACS patterns in annual hours worked across years are similar for most demographic groups. There are two exceptions. First, similar to what others have documented in the literature, annual hours works in the 2000 CPS exceed hours worked in the 2000 Census. (See, for example, Clark et al. (2003).) Despite the differences in levels of hours worked, the relative changes in annual hours across sex-age-education groups are very similar. As in the CPS data, less educated younger men had the largest decline in annual hours during the 2000s, decreasing by 63 hours per year, more than for less educated older men from 2000 to 2015. This pattern is nearly identical that shown in Table A3 based on the CPS. Also similar to Table A3, younger and older men with 4-year degrees had nearly the same decline in annual hours during the 2000s. The second difference to note is that the Census/ACS data show only a small trend difference in hours worked between younger and older women with 4-year degrees. This difference was much larger in the CPS data.

Panel (b) of Table A4 explores robustness to excluding full-time students ages 25 and older. The patterns between Panel A and Panel B are nearly identical. Annual market hours for less-educated younger men declined by 172 hours when all full time students are excluded. The comparable number is a decline of 183 hours per year when only full-time students under age 25 are excluded. This suggest that our ATUS and CPS results are not substantively affected by including full-time ages 25 and older.

A5 PSID Sample and Consumption Measures

To analyze the potential insurance younger men receive from parental and government payments, we examine the importance of transfer receipts for households in the PSID data that include younger men. From the PSID, we also examine these households' expenditures on non-durables and services. Results are discussed in the text. Here we describe our PSID sample and provide further description of how we measure consumption in the PSID.

We use 2001 to 2013 biannual PSID surveys. Our data primarily derive from the PSID Family Files, which contain information on income, transfers, and expenditures. We augment these with data on indivdual household member characteristics from the PSID cross-year Individual Files. We exclude the SEO and Latino special samples. Households are weighted by the Longitudinal Core/Immigrant Family Weight. Our sample size, in terms of household years, is 6,634 for households men ages 21 to 30; it is 16,155 for the reference group of households with men ages 31 to 55. (Note these two groups of households are not distinct sets.) Given the framing of PSID questions, reported prior year income and expenditures in survey year t are associated with calendar year $t-1$.

Table A4: Annual Hours Worked During the 2000s By Age-Sex-Education Groups, ACS

(a) Excludes Full Time Students Under Age 25

	Men Ed<16		Men Ed≥16		Women Ed<16		Women Ed≥16	
	21-30	31-55	21-30	31-55	21-30	31-55	21-30	31-55
2000	1,749	1,884	1,937	2,197	1,231	1,314	1,630	1,560
2007	1,712	1,849	1,913	2,169	1,196	1,309	1,638	1,563
2010	1,478	1,665	1,817	2,109	1,116	1,248	1,624	1,579
2015	1,567	1,764	1,859	2,125	1,176	1,253	1,663	1,630
Δ 2000-15	-183	-120	-78	-73	-55	-61	33	70
%Δ 2000-15	-11.0%	-6.6%	-4.1%	-3.4%	-4.5%	-4.8%	2.0%	4.4%

(b) Excludes All Full Time Students

	Men Ed<16		Men Ed≥16		Women Ed<16		Women Ed≥16	
	21-30	31-55	21-30	31-55	21-30	31-55	21-30	31-55
2000	1,760	1,888	2,013	2,216	1,230	1,314	1,665	1,562
2007	1,732	1,855	2,002	2,189	1,195	1,310	1,692	1,564
2010	1,499	1,675	1,916	2,129	1,116	1,253	1,681	1,585
2015	1,589	1,770	1,950	2,142	1,178	1,254	1,722	1,634
Δ 2000-15	-172	-118	-63	-74	-52	-60	57	72
%Δ 2000-15	-10.3%	-6.5%	-3.2%	-3.4%	-4.3%	-4.7%	3.4%	4.6%

Note: Table shows annual hours worked from the 2000, 2007, 2010, and 2015 ACS. Annual hours equal weeks worked over the last 12 months multiplied by usual hours worked per week. (ACS respondents report weeks and usula hours worked during the prior 12 months.) Panel (a) excludes full-time students ages 24 and under. Panel (b) excludes full-time students regardless of age.

The PSID provides data on non-durable and service expenditures at the household level, while our analysis on employment and hours concerns individuals. We take a standard approach by deflating household expenditures by a measure of household scale (equivalence units). We set this scale equal to $\sqrt{n}$, where n denotes number of household members. A square-root scaling factor is adopted in recent OECD studies (www.oecd.org/social/inequality.htm). Note that we treat all household members symmetrically. Thus, in a household with a working prime-age adult plus a non-employed younger man, we would allocate an equal amount of consumption to both. To the extent that the expenditure of such households are geared towards the parents, we will overestimate consumption of these younger men.

In Table A5 we report the growth rate in average expenditure for all households that include younger men ages 21 to 30. For comparison, we report the same for households that include men ages 31 to 55. These sets overlap to the extent younger and older men

Table A5: Real Consumption and Income Growth from 2000 to 2012, PSID

	Men: All Ed		Men: Ed<16	
	After-tax Income Growth	Consumption Growth	After-tax Income Growth	Consumption Growth
Households w/ Men 21-30	-6.6%	-0.7%	-10.0%	-4.8%
Households w/ Men 31-55	-3.9%	-5.5%	-10.0%	-6.7%
Difference	-2.6ppt	4.8ppt	-0.04ppt	1.9ppt

Note: Data reflect 2001 and 2013 PSID surveys, corresponding to calendar years 2000 and 2012. Series are deflated by household-specific equivalence scale and the GDP deflator. The household equivalent scale equals the square root of number of household members. After-tax income is calculated by netting taxes from the before-tax income reported in the PSID, where taxes are calculated using NBER TAXSIM. The consumption measure reflects expenditures reported on rent, or imputed rental equivalence for owners, utilities, food, transportation (gasoline, public transit), heath, and education.

are coresidents. Our measure of consumption includes expenditures on housing (either rent or imputed rental equivalence for owners, and utilities), food (both consumed at home and away), transportation (gasoline, public transit), heath, and education. These are the NIPA-defined nondurable and service categories reported consistently across the 2001-2013 PSID samples. (Rental equivalence is imputed based on owner's reported home value. This mapping is estimated from the BLS Consumer Expenditure Survey, which contains responses on rental equivalence and on home value.) The table also reports the growth in household after-tax income for each subgroup. Before-tax income reflects PSID responses, while household taxes are calculated using NBER TAXSIM. Both income and expenditures are deflated by each household's equivalence scale, discussed above, and the GDP deflator.

Looking at the first two column's of Table A5, we see that households with younger men displayed only a slight decline in real expenditure, 0.7 percent, despite displaying a decline in household income of 6.6 percent. The table compares results for younger and older men. We see that households with younger men displayed a 4.8 higher growth in consumption than households with older men, despite displaying 2.6 percent lower growth in income. The latter columns of Table A5 again compares growth in expenditures for households with younger men versus older men, but now restricting attention to men with less than four years of college. Again we see slightly higher growth in expenditures for the younger men, by 1.9 percent, while household income growth looks the same across the two groups.