Labor Supply and Entertainment Innovations: Evidence From the U.S. TV Rollout*

George Fenton and Felix Koenig†

June 10, 2024

Abstract

We study the impact of entertainment technologies on labor supply using a natural experiment afforded by the regulated U.S. television rollout. We find that TV significantly affected retirement rates but had little impact on labor supply among prime-age workers. The launch of a TV station reduced the probability of working by around 0.3 percentage points, driven mainly by an increase in retirement rates among older age groups. The results thus support the notion that the rise of television contributed to the mid-century transition of retirement as a necessity to the enjoyment of old age. However, our findings suggest that entertainment innovations have a less pronounced effect on overall labor supply trends than model calibrations in the previous literature suggested.

*We are grateful to Ben Olken and Matthew Gentzkow for sharing code and data, Hannah Rhodenhiser and Alex Cheng for diligent research assistance, and to Andreas Ferrara, Brian Kovak, Adrian Nieto, Steve Pischke, Mel Stephens, Lowell Taylor, Justin Wolfers, and seminar participants at Georgetown, Michigan, LSE, MIT, University of Pittsburgh, Princeton, and Lausanne for comments. This research was supported by the Michigan Institute for Teaching and Research in Economics and by an NICHD training grant to the Population Studies Center at the University of Michigan.

†Fenton: U.S. Department of Treasury; Koenig: Carnegie Mellon University (fkoenig@cmu.edu)

> *Here we must begin with the most fundamental fact about the impact of television on Americans: Nothing else in the twentieth century so rapidly and profoundly affected our leisure.*
>
> Robert Putnam[1]

> *The fact that even in 1950 the average television household was watching for four and a half hours per day makes clear what a dramatic improvement television was over previous entertainment technologies.*
>
> Matthew Gentzkow[2]

1 Introduction

In a classic labor supply model, whether to work and how many hours to work depend on the relative value of work and non-work time. A large literature examines the role of wages in determining labor supply, but less attention has been paid to changes in the value of non-work activities and their impact on work decisions.

The idea that the returns to non-work activities affect labor supply decisions was formalized in groundbreaking work by Becker (1965). More recently, studies took an interest in the impact of entertainment technologies and discuss how such technologies may affect the value of leisure time. Aguiar et al. (2021) find that video games led to an increase in the value of leisure for young men and a sharp decline in their labor supply. Similarly, Costa (1998) suggests that more varied and affordable leisure activities contributed to the rise of a "golden age of retirement" in the mid-twentieth century. Over the past century, home entertainment technologies evolved rapidly and new innovations have made entertainment technologies increasingly compelling and readily available. Today the vast majority of free time is spent using television, computers, and similar technologies, however, there is still limited well-identified evidence on the impact of such technologies on the labor-leisure trade-off and labor supply.

In this paper we study television, which quickly became Americans' dominant leisure activity, taking up more time than any other activity except sleep and work. Already during the early days of television, Americans spent on average more than ten hours per week watching TV (see Figure 1), and the most active percentile spend as much as 40 hours per week in front of the TV.[3] This shift

[1] See Putnam (1995), p. 221.

[2] See Gentzkow (2006), p. 970.

[3] This data only accounts for television watching as a "primary activity." There is additional time spent with television as a secondary activity, i.e. watching television while performing other activities.

towards watching television represents one of the biggest changes in American time use over the last century, arguably making it the most significant entertainment innovation of the modern era.[4] Because of TV's ubiquity today, identification can be a challenge. A review of related literature by Abraham and Kearney (2020), for example, concludes, "We do not attempt to assign a magnitude to the possible contribution of improved leisure technology... This is an issue deserving additional attention" (p. 52).

The main empirical challenge with studying the impact of television is that individuals with large amounts of spare time self select into television viewing. We use two natural experiments to overcome this identification challenge. The first strategy leverages the fact that television broadcast towers were deployed by the government in a staggered fashion, which generates variation in the timing of television's introduction across local areas (a design pioneered in Gentzkow (2006) to estimate TV's effect on voter turnout). The second strategy uses novel data on hold ups in the government deployment process that arose for regulatory reasons. These unexpected delays in the rollout generate "ghost stations" that were meant to go live but could not because of the interruption. The interruption started in September 1948 and was expected to last about six months, but was ultimately not lifted until nearly four years later, creating credible treatment and control groups during this period. Prior work proxied for locations affected by the freeze with stations that launched shortly after the freeze was lifted. This is an imperfect approximation because the priority rankings were revised during the interruption and different rules were used to select stations before and after the interruption, potentially making places whose stations launch after the freeze less similar to the pre-interruption places than one might hope.

We use novel data to directly measure which areas where affected by the freeze. This allows us to run two types of tests that help with identification. First, we compare treated areas (where applications were approved) only to areas where applications were frozen, rather than to the entire untreated sample. Second, we conduct a placebo test that treats frozen stations *as though they had been approved and launched* and estimates the impacts of such "ghost stations." We find that the frozen stations have no effect on local labor supply, suggesting that station launches are not spuriously correlated with local economic conditions, adding credibility to the baseline DiD results.

The baseline DiD specifications show statistically significant but modestly sized impacts on work. Specifically, we find that television had a meaningful impact on the retirement decisions of older workers and limited effects on the labor supply of prime-age workers. DiD regressions of individual Social Security employment records on TV exposure show that the launch of an additional channel is associated with a decline in the probability of working on the order of 0.3 percentage

[4] In a study of time use in the twentieth century, Aguiar and Hurst (2007) find that "More than 100 percent of the increase in leisure can be accounted for by the increase in the time spent watching television" (p. 987).

points. The impact on the entry and exit rates of workers under the age of 40 are statistically and economically insignificant, while the exit rate for workers over 50 increases significantly. These exits are predominantly into retirement. Retirement rates among older workers increase by 0.25 percentage points, suggesting that on average people retire 2 months earlier. We show that these results are consistent with the economic models of retirement and the intuition that people at the margin of labor force participation are most likely to respond to changes in the value of leisure. They also support Costa (1998)'s hypothesis that the greater availability of compelling, low-cost entertainment like TV contributed to the large mid-century increase in retirement and changed the perception of retirement from a mere necessity to "golden years" of relaxation.

We probe the validity of our findings with further robustness checks beyond the "ghost station" quasi-experiments outlined above. Our preferred specifications rule out that the results are driven by composition changes in the sample by including individual fixed effects in the regression. We next test and confirm that differential trends in labor force participation across demographic groups that could have correlated with TV's rollout also cannot account for our findings. Similarly, pre-trend tests show no evidence of spurious pre-treatment trends. We also consider how migration across local labor markets may affect our results. The baseline specification keeps individuals' locations fixed, which helps rule out confounding effects but also introduces measurement error in television exposure, possibly attenuating our results. We provide two robustness checks to assess this potential issue. We first run the baseline regressions for a sample of individuals who are less likely to have moved and second test directly if the timing of television is correlated with local migration rates. The results suggest that migration is orthogonal to the TV rollout. We also use a bounding exercise to assess the magnitude of the potential problem and find relatively minor effects for plausible migration patterns.

Finally, we study the implications of our findings for the impact of leisure innovations on labor supply more broadly. Home entertainment has undergone a massive expansion in variety, quality, and availability over the past century, from the early advent of radio and TV to more recent innovations like YouTube and Netflix. We build a representative agent framework similar to Aguiar et al. (2021) to study the likely impact of innovations beyond our empirical context. One challenge in comparing alternative entertainment technologies is that the magnitude of such shocks is typically unobserved: technology shocks have no natural units and their impact on the value of leisure is therefore hard to quantify.[5] We develop a revealed preferences approach to quantify the impact of technology shocks on the value of leisure. The approach uses the allocation of time *within* leisure across alternative activities; a large shift in time use towards one activity implies a large improvement in the value of this activity. Since time shares are observable, we

[5]Some studies use "hedonic" price indexes to quantify product quality improvements. Ultimately, the credibility of these measures rests on correct proxies for product quality.

can link our empirical estimates to relevant preference parameters without directly quantifying the size of technology shocks. This approach is in the spirit of a "sufficient statistics approach" and derives simple reduced form results that identify the structural preference parameters of the model. The chief advantage of such an approach is that it enables researchers to combine the empirical credibility of quasi-experimental tools with the flexibility of structural approaches to study important policy questions beyond the specific empirical context.

The paper then uses the theoretical framework together with the empirical results to quantify how much labor supply is affected by television and other subsequent entertainment innovation. Relative to calibrations used in the literature our estimates suggest a smaller role for entertainment technologies like TV. The key preference parameters that determine results are substitute elasticities between alternative leisure activities and between leisure and labor. Our results imply that alternative leisure activities are much closer substitutes for each other than for labor (or consumption). Such preferences imply that workers are reluctant to reduce labor supply in response to improvements in the value of leisure and only the most major entertainment innovations (like television) will shift the value of leisure enough to yield economically meaningful effects. If preferences are similar in other settings, we can compare the calibration choices to our estimates. Our result implies that standard calibrations likely overstated the impact of entertainment technologies by assuming too small substitution elasticities between alternative leisure activities.

Our study contributes to three broad research areas. The first is on secular employment and retirement trends (for reviews see Abraham and Kearney (2020); Kopecy (2011); Vandenbroucke (2009); Greenwood and Vandenbroucke (2008); Juhn and Potter (2006) and Lumsdaine and Mitchell (1999)). The decline in participation rates among the elderly in the middle of the twentieth century was one of the largest shifts in U.S. employment rates over the past century (Blundell et al. (2016); Lumsdaine and Mitchell (1999); Costa (1998)). Costa (1998) suggests that "the lower price and increased variety of recreational goods has made retirement more attractive" and fostered a new "retirement lifestyle." Several models of retirement trends support this claim (Kopecy (2011); Vandenbroucke (2009); Greenwood and Vandenbroucke (2008)). We show both theoretically and empirically that, while TV affected everyone, the biggest responses occur at the retirement margin. Our study thus provides direct empirical evidence of the "retirement lifestyle" channel.

The second broad literature studies how home technologies influence labor markets. The idea that the value of leisure is an important determinate of labor supply goes back at least to Becker (1965). In "A Theory of the Allocation of Time," Becker argues that labor supply research primarily focuses on the opportunity cost from foregone earnings but is "not equally sophisticated about other non-working uses of time."[6] Since then, research on dishwashers, microwaves, washers, and

[6]"Most economists have now fully grasped the importance of forgone earnings, [...] it is perhaps surprising that economists have not been equally sophisticated about other non-working uses of time." (Becker, 1965)

dryers has found that such appliances acted as "engines of liberation" and increased women's labor force participation by reducing the burden of home production (See Greenwood et al. (2005) and related work by De Cavalcanti and Tavares (2008), Coen-Pirani et al. (2010), Ngai and Petrongolo (2017), Greenwood et al. (2016), and Bose et al. (2020)). Nieto (2020) argues that the launch of digital TV in the U.K. between 2008 to 2012 similarly acted as a substitute for child care, which increased women's employment.

Studies specifically on the impact of technologies on the value of leisure are comparatively scarce. Several papers model the impact of entertainment technologies on macroeconomic trends and suggest that such technologies may have played a major role. Most relevant, Aguiar et al. (2021) examine how video games affected the labor supply of young men during the 2000's. Similarly, Kopytov et al. (2023) and Rachel (2020) find that declining prices of entertainment technologies could rationalize broader employment trends, and González-Chapela (2007) finds a negative correlation between local entertainment prices and labor supply. Some scholars flag the absence of clean identification as a challenge in these settings. Abraham and Kearney (2020) note, "the mechanism and direction of the effect warrant consideration, but the point estimates reported unavoidably rest on a good many unverifiable modeling assumptions." Our study leverages a natural experiment to provide a well-identified estimate of the impacts of leisure technologies.

Finally, our work also relates to a growing literature on the role of television in society (including work by Gentzkow (2006); Gentzkow and Shapiro (2008); Jensen and Oster (2009); Baker and George (2010); Campante and Hojman (2013); Kearney and Levine (2015); Mastrorocco and Minale (2018) ; Thomas (2019); Kearney and Levine (2019); Nieto (2019); Kim (2020); La Ferrara et al. (2012); Angelucci et al. (2021); Chadi and Hoffmann (2021) and Gennaro and Ash (2022)).

The rest of this paper is organized as follows. In section 2, we present our conceptual framework and derive testable predictions; in section 3, we discuss the data; section 4 presents the design and main results; section 5 discusses the implications of the findings; and section 6 concludes.

2 Data

Our study combines a newly built data set on television signal strength in the 1940's and 1950's with administrative employment records.

2.1 Measuring TV Access

To date, there are no comprehensive measurements of TV signal strength during the U.S. rollout. Previous studies typically approximate the coverage of 1950's stations with the boundaries of Designated Market Areas (DMA's) from the 2000's.[7] We digitize archival records to precisely measure television signal reach. The chief advantages of the new data set are twofold. First, we more accurately measure the broadcast boundaries of each given station; and second, we measure coverage intensity–the *number* of stations available in an area–which makes for an improvement over the binary DMA approximation of TV availability.

Commercial television was first licensed for broadcast in 1941, with experimental stations in a few major cities like New York and Los Angeles. The rollout took off after World War Two, and the post-war expansion was a staggered city-by-city process over the following two decades whose timing was governed in part by a sharp regulatory freeze. The freeze came about due to signal interference between neighboring stations, an issue that occured due to an error in the Federal Communications Commission's (FCC) signal model. No new stations were allowed to launch until a new system was implemented that would ensure interference free transmission. This interruption in the roll-out plays an important role in our identification strategy and we return to this topic below. Most of the growth in coverage and viewership occurred in subsequent years, during the 1950's; in 1950, less than 20 percent of households owned a TV, and by 1960, 87 percent did (see Gentzkow (2006) for a detailed discussion of the rollout process). Our first contribution is to produce precise measurements of TV access in this period.

We use the Irregular Terrain Model (ITM) to calculate signal reach during the rollout. The ITM computes signal strength in decibels at a receiving location as a function of the distance of that location from a broadcast tower, tower technical specifications, and topography between the tower and receiving location.[8] The new data has two advantages. First, we reduce measurement error and discuss such improvements in detail in Appendix A. Second, the DMA approximation ultimately produces a binary coverage variable. Since different cities also had different numbers of stations, and some pioneering stations had limited broadcast hours, a binary treatment indicator can miss variation of interest in the intensity of TV "treatment." With the ITM, we can separately calculate signal strength for each individual station and therefore track the rising availability of TV at both the extensive and intensive margins.

Using the ITM requires detailed information on broadcast towers. We collect three sets of data on broadcasting technology from early editions of the *Television Factbook*, a trade publication for advertisers and other industry players. First, beginning in 1948, the *Factbook* published the technical characteristics of all commercial stations in operation. We use these as inputs for the ITM. Specifically, for each station in each year from 1948 to 1960, our digitized *Factbook* data

[7]Work using this DMA approximation includes: Gentzkow (2006); Gentzkow and Shapiro (2008); Baker and George (2010); Campante and Hojman (2013); Thomas (2019); Kim (2020); and Angelucci et al. (2021)

[8]The ITM model has also been used in other countries: Olken (2009); Enikolopov et al. (2011); Della Vigna et al. (2014); Yanagizawa-Drott (2014); and Durante et al. (2019). Wang (2020) also uses the ITM to estimate the effects of a 1930's populist radio program in the U.S..

include latitude and longitude, height above ground, channel number and frequency, visual and aural power, and other details like call letters and start date. There were 41 stations on air in 1948. Already by 1960, there were 570.[9] We estimate the signal strength of each station at the geographic center of each U.S. county from 1948 to 1960.

The second and third groups of data involve secondary extensions of original broadcasts. A town across a mountain range from a nearby city would be cut off from that city's TV signals, and the ITM would correctly measure that town as having no TV access through the air. However, some towns constructed antennas on top of the mountains to capture signals and then wire the broadcasts into the otherwise obstructed homes. This was the birth of cable TV and was known at the time as Community Antenna Television (CATV).[10] We have digitized the *Factbook* directories of CATV locations, start dates, and estimated number of subscribers. Finally, an alternative to piping a signal through a CATV system was to rebroadcast it through the air with small antennas called translators. The *Factbooks* record the locations of licensed translators beginning in 1957, and we have digitized them through 1960.

Figure 2 shows a snapshot of the ITM output in 1950. Here we have mapped the strongest signal available in each county. The units are decibels, where zero indicates top-quality signal strength. Any signal below -50 decibels was effectively unwatchable, and we have colored the figure to indicate that coverage transition as the map shifts from red to blue. City centers are clearly visible, as is the fading strength of the signals–a typical broadcast reached about 100 miles from its tower, leaving some counties well outside of urban centers still within reception rings but others out of range. This is an extensive margin perspective on the data, in the sense that the map displays whether a county could receive a watchable signal from any station. We also estimated the number of stations available in each county in each year. Summary statistics for the time path of the roll-out are presented in Appendix Figure 9.

The timing of launches will play an important role in our identification strategy. Launch decisions are clearly not taken at random and we devise several strategies to isolate policy idiosyncrasies for identification purposes. An unexpected FCC licensing freeze halted approval of all new stations from September 1948 to April 1952. Stations whose applications were approved before the 1948 freeze were allowed to continue broadcasts, but those pending approval when the freeze took place could not begin broadcasting until the freeze was lifted four years later. The sizable impact of the end of the rollout interruption can be seen in Appendix Figure 9 in a jump in television launches. We use additional data on frozen applications from Koenig (2023) and combine it

[9] Latitude and longitude are first published in the 1952 *Factbook*. Earlier years give station addresses, which we geocode. The *Factbook* was published four times per year in 1948 and 1949 and twice per year from 1950 to 1960. We digitize the latest edition available in each year.

[10] In 1966, both the *American Economic Review* and the *Quarterly Journal of Economics* published articles on CATV; see Fisher (1966) in the references.

with the ITM model to implement a novel empirical strategy that computes the signal strength of stations that were in reality blocked by the FCC. We use this data for a powerful placebo test that treats these stations *as though they had been approved*. If a regression specification using these "ghost towers" shows effects of TV where there was none, then that specification must reflect spurious correlations. Reassuringly, we find no effects from "ghost towers."

2.2 Employment Data

Our main source of labor market data is the Current Population Survey Social Security Earnings Records Exact Match file (henceforth "SSA-CPS"), which matched respondents from the March 1978 CPS to their entire Social Security earnings histories going back to the 1930s.[11]

The data is a worker-level panel and is one of the only micro data sets that covers years between the decadal Censuses during this period. We focus on the adult population (aged 21+ at the time) in the mainland U.S. and study changes in working behavior between 1937 and 1960. The sample includes 292,448 worker-year observations, of which people close to retirement will turn out to be particularly relevant (48,075 observations are aged 50+ in our sample window). Appendix section 10 provides summary statistics and further data details.

A major appeal of the data is the panel feature, which allows us to track the same individuals over time. We can thus run the analysis at the micro-level and use individual fixed effects to hold individual characteristics constant. A further appeal of the administrative data is that the records are based on employer reports to the SSA, which tend to be more accurate than retrospective survey questionnaires. A drawback of administrative data is usually the lack of detailed demographic information. Since our data is based on the CPS, we can link the SSA records to information from the CPS. This allows us to use information on workers' age, race, education, occupation, and place of residence. The residence information is the metropolitan statistical area (MSAs) and rest of state for non-MSA residents, and we run the regressions at this geographic level. Statistical agencies expanded the granularity of spatial information over time, we use the areas identified when our data was published and use the same 134 locations for all sample years.

For each individual we observe the number of qualifying quarters worked per year and we code an individual as employed if they work roughly half a quarter. Specifically, if their earnings exceed the level required for half a qualifying quarter.[12] The data reports are at annual frequency during the 1950's, however, in earlier years multiple years are grouped together and multi-year summary

[11]This dataset was initially compiled by the Bureau of Labor Statistics to evaluate survey responses in the CPS; aside from such evaluations, the data has been underutilized by researchers. A notable exception is Acemoglu et al. (2004) who study labor supply behavior of women in the post-war period. The data is available as ICPSR repository 9039.

[12]SSA qualifying quarters may differ from quarters worked if earnings in a quarter are below the qualifying threshold or if a person works in non-qualifying employment (e.g. some self-employment).

records are available for 1937-1946 and 1947-1950. The baseline sample includes the annual data for 1951-1960 and one observation for each of the multi-year bins, so there is one observation for 1947-1950 and one for 1937-1946. In these multi-year observations we take variable averages for the period. Since averaging over multiple years inevitably introduces noise, we down-weight the multiyear bins by the inverse of the number of years in the bin (using weights of $\frac{1}{10}$ and $\frac{1}{4}$ for the two bins, respectively).

We additionally account for the expansion of Social Security coverage and the Korean War in the 1950's. The Social Security administration expanded their definition of employment during the 1950's. We drop individuals who are affected by the coverage expansion to work with a consistent sample (See Appendix section 10 for further details). The start of the Korean War led to a draft and we exclude drafted soldiers from the analysis to avoid spurious employment effects from the draft.

There are two potential challenges with this data. The first is that our demographic information is only collected in 1978 in the CPS. Importantly this means that we observe the place of residence only in 1978. We follow previous work by Acemoglu et al. (2004) and treat demographic information as fixed throughout the sample period. Since this may introduce measurement error in our TV exposure variable we provide several robustness checks and bounding exercises to assess the impact of this assumption. Specifically, we find that the timing of TV launches is uncorrelated with the local share of out-of-state retired residents and our main results remain similar if we restrict the sample to people who are more likely to reside in the same location throughout the sample (see (Kearney and Levine, 2015) for a related challenge and approach).[13] Still one may worry about attenuation bias and we use a bounding exercise to assess their potential magnitude and find that the impact on the results is small for realistic migration patterns (see Appendix 10.3.4). A second data challenge is that the 1978 CPS cohort is not representative of the U.S. in earlier years. Specifically, we will observe fewer people in cells with lower survival rates. For example, while our sample includes over 45,000 observations aged over 50, the sample is thin on individuals over 70 year old.[14] Studying the impacts of TV on a population that is younger than average is not a problem for the internal validity of our results. However, one would want an estimate that is representative for the entire population if one is interested in comparing our estimates to macro trends. To provide such a more representative estimate, we re-weight our sample to account for different survival rates and find similar results to the baseline, with slightly bigger effects (Appendix 10).

We use an additional data source to study hours worked and intensive-margin labor supply re-

[13]Different from Kearney and Levine (2015), our results are not affected by potential migration induced changes in the composition of local labor markets since we use individual panel data and can hold individual characteristics constant.

[14]Arguably, individuals over 70 are uninteresting for a study of labor supply at mid century, since the vast majority of them were economically inactive.

sponses. In the 1950s, data on work hours were collected for national statistics but rarely reported at geographically disaggregated levels. An exception are several Bureau of Labor Statistics publications that provide local area breakdowns of hours data. These records are summarized in the Current Employment Statistics (CES) and report average hours worked among non-agricultural employers in the manufacturing sector across different local areas. The reporting areas are typically MSAs or state level aggregates. Ideally, we would use the same geographic areas as the baseline analysis, but the available data mostly covers areas not identified in our baseline data. The hours data is published for 51 local areas during the period 1947-1960.[15] The panel is thus relatively small but provides a glimpse into intensive margin effects.

3 Empirical Analysis

We now study the impact of television on labor supply. Television station launches were staggered over two decades, leading to substantial regional heterogeneity in access. As a first pass, we use all station launches, whether affected by the rollout interruption or not and run the following difference-in-differences regression:

$$\Delta E_{aigt} = \gamma_{gt} + \delta_i + \beta_k \cdot TV_{at} + \pi \cdot X_{aigt} + \epsilon_{aigt}, \qquad (1)$$

the outcome variable ΔE_{aigt} measures changes in employment and E_{aigt} is an employment indicator with value 100 if individual i of gender g in area a at time t is employed. TV_{at} measures the number of available TV stations in area a at time t. The baseline specification estimates the average effect of an additional station and we explore diminishing effects and other heterogeneity further below. The baseline specification does not restrict attention to the first stations because many of these first launches happen before our outcome data become annual. Time-fixed effects (γ_{tg}) absorb aggregate trends in labor supply (e.g. driven by policies and changing norms) and nationwide general equilibrium effects. We allow for different year effects by gender (g) since employment trends were different in the post-war period. Individual fixed effects (δ_i) control for individual characteristics and cohort specific work patterns; these also absorb area effects, since we assign individuals to a time-invariant area a. Finally, X_{aigt} is a vector of control variables. The effect of TV on employment is captured by β_k and we allow this effect to differ between demographic groups denoted by k. Specifically, we will allow for different impacts on younger and older workers and on men and women.

This baseline DiD requires that treatment and control areas are on parallel trends, and the disruption experiment will help us relax this assumption below. We also provide a battery of

[15] Coverage of local areas expands throughout this period and the data is therefore unbalanced.

robustness checks to investigate whether parallel trends hold in this DiD set-up.

3.1 Baseline DiD Results

We first present event-study graphs that plot changes in employment before and after the launch of a local TV station (Figure 3). We show these patterns separately for older workers in panel A (aged 50+) and younger workers in panel B (under 50 years old). Both Panels show that treatment and control regions evolve in parallel in the years leading up to the launch of a TV station. After the launch of a TV station, we see a sizable decline in employment among older workers in affected areas, whereas there is little change in work patterns among younger individuals. Among older workers TV reduces employment around 0.3 percent. The effect on older workers sets in the year TV is launched and consistent with the notion that retirement takes some time to adjust, the effect increases over subsequent years. The cumulative effects of TV thus increase over time, while the results do not show signs of spurious pre-trends. The point estimate for younger workers is roughly a third of this size of the effect on older workers.

We next turn to standard DiD regressions, following equation (1). This standard DiD pools data winthin pre and post periods and thus boosts the power relative to the even-study period by period estimates. Table (1) shows that an additional TV station reduces employment among people over 50 years of age by 0.3%, a significant negative effect, while the effects on younger workers are small and insignificant. The point estimate for the younger group is around 0.1 percent, indicating that prime-aged individuals responded much less to the launch of television. The labor force participation rate in 1940 was around 63% for the younger group and 44% for the older age group.[16] The changes are thus moderate in size and, in the case of prime-aged workers, indistinguishable from zero.

The impact on older workers is more economically interesting and helps explain the change in retirement trends taking place at the time. Below, we will show that a key driver of the results are changes in retirement among older workers. Suppose the effect was entirely driven by retirements then the 0.2-0.4 percentage point change in life time employment is equivalent to a reduction in retirement age of around one to four months (for a life expectancy of 71, the effect is: $0.002 \times 71 \times 12 = 1.7$ months). We also investigate whether women responded differently to the introduction of television. On the one hand, it seems plausible that weaker norms about labor force participation make women more responsive. On the other hand, if major labor supply decisions, such as retirement, are taken jointly at the household level, one might expect similar behavior for men and women (evidence for joint retirement is presented in Gelber (2014)). In practice, we find that women respond similarly to men and a coefficient that allows for different effects by gender

[16] Source 1940 Census, Tabulation uses IPUMS tabulation tool created by CSM, UC Berkeley

is statistically insignificant.

Next, we check for the robustness of these results to several identification threats. We first address the fact that employment trends between men and women diverged sharply in the second half of the twentieth century. Our regression control for the impact of such gender-specific trends with gender-specific year effects. A second potential concern are changes in the composition of the local workforce. This issue is prominent in studies of local labor market level outcomes but is not an issue in this study. The regressions are run at the individual level and individual fixed effects ensure that composition changes do not affect the results. Our preferred specifications additionally control for age group fixed effects to capture that employment probabilities evolve over the life cycle. The variation that identifies these preferred estimates compares changes in employment status among workers of the same age group in places that gain TV access versus once that don't. Results with these additional controls are similar to the baseline (from Column 2 onwards).

The main identification assumption is parallel trends in treatment and control areas. We now provide additional checks to probe this assumption beyond the event-study graphs presented above. Previous work indicates that richer and more densely populated areas received TV before other parts of the country (e.g., Gentzkow (2006)). The potential threat is about the resulting trends from imbalanced demographics; while effects from imbalances in local characteristics themselves are absorbed by location-fixed effects. The institutional features of the rollout give us reason to be optimistic the parallel trend assumption holds. The FCC processed launch permits according to its internal priority ranking of locations. The position in this ranking was based on largely fixed location characteristics (e.g. in 1956 on population and distance to nearest antenna). An important implication of this is that changes in local conditions did not have an effect on the timing of television launches.

To detect potential confounding trends, we implement a balance test style analysis that checks which demographic trends predict TV launches. Specifically, we run a local area level regression with the TV station variable as outcome and regress it on local area demographics (city status, gender ratio, average age, high-school completion, race, marital status, employment rate, average income, four indicators for Census regions) and interact the local demographics with a linear time trend, controlling for year and location-fixed effects. We find only a limited predictive power of local demographic trends. Including all 11 demographic trends in the regression increases the regression R^2 by 0.017, suggesting that our identification variation is largely exogenous to these trends. The coefficients are plotted in Panle A of Figure 4 and are insignificant in all but two cases.[17] The significant predictors are indicators for two Census regions (South and North), where TV was wide spread earlier than in the omitted region (the western part of the country).

[17]Interestingly, the significant variables are different to the ones significant in a similar balance test in Gentzkow (2006), reflecting the differences in the specifications and TV variation in the two studies.

To address the potential identification threat from these region-specific trends, we next add region-specific time trends as control variables to the DiD regression. A first specification interacts a fourth-order time polynominal with the four indicators for Census regions and finds results close to the baseline estimates (Column 3). We then replace this parametric time trend with fully flexible time effects and interact Census regions with year effects. The identification variation in these specifications comes from the variation within census regions and the results again remain similar to the baseline estimates (Column 4). Taken together, the results indicate that we identify effects that are orthogonal to region-specific trends.

A final specifications expands the set of demographic controls. We now include demographic trends for variables that are insignificant in the balance test but have sizable point estimates (bigger then 0.005 in absolute terms). These specifications include the previous region trends and add trends by marital status and baseline local employment rates. We again allow for a quadratic trend in these. The results are similar to the baseline (Columns 5).

The next part of the analysis explores how the effect of stations diminishes with a growing number of available stations. The baseline results used a linear regression in TV stations and thus estimates an average effect of all launches. We now allow for non-linear effects. A first specification allows for diminishing effects by introducing a quadratic term of the TV regressor. The results confirm that the effect of stations wears off once many stations are active in an area. The negative point estimate on the linear TV variable is partially offset by a positive coefficient on the square term (see Table 2). Both coefficients are highly significant, suggesting that the main drivers of the effects come from the initial stations and the effect of an extra station is smaller after a few station launched (Column 2). We next allow for an alternative functional form and replace the TV station regressor with log TV stations (using ln(1+station count)), imposing a logarithmically decreasing effect of TV launches. In effect, this transformation reduces the influence of higher TV count values and puts greater weight on the first stations. The estimates show a similar pattern to our baseline estimates and confirm again that the initial stations are the main driver of the results (column 3). Both these approaches impose specific functional form assumptions, and we next relax this assumption and explore less parametric estimates. Specifically, we run specifications that top code the station count variable. These specifications focus only on the first X station launches. The first specification narrows in on the intitial station and top-codes the station count at one station (Column 4). Note that this is not quite the same as a binary yes/no TV variable. Our station count variable measures the average number of stations and if some but not all individuals have access to signal, this can create non-discrete TV counts, making the top coded variable close to but not the same as a yes/no TV variable. A second specification top codes the station variable at five stations (Column 5). Qualitatively, these specifications show the same pattern as our initial baseline estimates above: TV has negative effects on employment, mainly coming from changes among older workers. The magnitude of these estimates declines the more stations we take into account. We find the largest effect when focusing on the first station only (point estimate of -0.9), a somewhat smaller effect when estimating the average effect of the five first stations (point estimate of -0.6) and again smaller effects when estimating the average over all launches (baseline estimate of -0.3). All in all, we find a consistent pattern of larger effects from the first stations and smaller additional effects when multiple stations are already available.

We next address the impact of migration on our estimates. First, recall that our analysis treats individuals location as fixed. This alleviates concerns that moves lead to spurious variation in access to television: for example, with locations kept fixed, a person who retires to Florida does not experience an endogenous change in their TV exposure at the time of retirement.[18] We thus rule out that relocation decisions generate spurious variation in TV access.

Treating individuals' residence as fixed may, in turn, introduce measurement error in our TV exposure measure. People who moved likely experienced TV launches at different times than we observe and such measurement error will attenuate our results. However, we show in Appendix 10.3.4 that the magnitude of such effects is small and in a bounding exercise show that the potential impact is minor for realistic migration patterns.

3.2 The Rollout Interruption Experiment

We now leverage the natural experiment created by the rollout interruption to provide further credibility to the results and probe the parallel trend assumption of the baseline DiD.

As a first step, we investigate whether places affected by the interruption look similar in terms of demographic characteristics to places where launches occured shortly before. We again implement this with a balance regression, regressing an indicator of being held up by the interruption on local demographics, and plot the estimates in Panel B of Figure 4. Characteristics look broadly balanced. The balance in the interruption experiment is slightly better than in the baseline DiD sample, although our power is reduced by the smaller sample. There are no statistically significant differences between treatment and control locations and the point estimates are closer to zero. Whether places were narrowly affected by the interruption looks as good as random.

We then use the interruption sample for several robustness tests. First, we investigate parallel trends by studying if places with launches experience different effects compared to places with planned launches that are blocked. The attraction of this test is that we observe places that were meant to be treated in an untreated state of the world. This variation detects potential spurious effects at the time period of the supposed treatment and is thus an even stronger test for paral-

[18] Note the trend towards retirement in the sunbelt couldn't explain our findings in any case. TV did not arrive early in Florida, nor in the rest of the sunbelt. More generally, the local share of the out of state retired population is uncorrelated with early TV exposure.

lel trends than a conventional pre-trend check. We first implement this test as a horse race in a regression with a treatment variable for station launches and one for planned but blocked station launches. The results show that negative employment effects only arise from launched and not from blocked television stations (Table 3, Columns 1 and 2), which provides direct evidence that the rollout rules are unrelated to spurious local labor demand shocks.

Second, we use the interruption natural experiment for a distinct identification strategy. This strategy arguably uses cleaner variation than the baseline DiD, but has the drawback that it focuses on a more limited set of locations, reducing our power and potentially the external validity. We implement the test with two different control groups: The first analysis compares places that received TV shortly before the interruption to places that received TV shortly after the interruption (Columns 3 & 4),[19] and the second compares the treated places to places that at the time of the interruption were ranked next in the rollout list (Columns 5 & 6). The first strategy has been used in the literature on the US television rollout (see, e.g., Gentzkow (2006); Gentzkow and Shapiro (2008); Baker and George (2010); Campante and Hojman (2013); Thomas (2019); Kim (2020); La Ferrara et al. (2012); and Angelucci et al. (2021)), while the second strategy is novel and arguably a cleaner experiment. To make the second strategy possible, we use novel data on the priority list from Koenig (2023) and augment this data with information on the counterfactual signal of the stations that did not launch. The two strategies differ in the locations they use as control group. The rollout list was revised during the interruption period and some places leapfrogged in the ranking. As a result, the first strategy potentially conflates variation from the interruption and variation from the change in priority criteria, while the second strategy uses only variation from the interruption.

The results again show a negative effect of television similar to our baseline findings. The pre-post interruption comparison shows a 0.2 percentage point decline in employment (Column 3). This is similar to our baseline estimate and thus adds confidence in the validity of the baseline results. While the interruption ensures that treatment and control groups in this experiment are broadly similar, there may still exist small differences in baseline characteristics. We therefore control for the demographic composition and allow for group specific time trends. Given the differences are small to begin with, it is unsurprising that adding these controls has little effect on the results (Column 4). We next use the priority list approach and additionally narrow in on the years around the interruption for an even cleaner experiment. The optimal time window for the experiment depends on the time it takes for effects to unfold and the variation we want to exploit. The interruption creates two potential "experiments." One uses the timing of the start and another the timing of the end of the interruption. Both generate quasi-random variation in the timing of TV launches but in practice, the lack of annual outcome data in the pre-interruption period means that

[19]Recall that the interruption lasted from September 1948 to April 1952; The pre-post interruption sample focuses on areas with launches between 1947 and 1954.

we cannot utilize the start of the interruption and focus on the end of the interruption. The sample window for this experiment should include the interruption period and sufficient post-interruption periods to allow the interrupted stations to take their full effect. The interruption begins in 1948, so we would ideally start the sample window in 1948. However, this year is part of the 1947-1950 multi-year bin and we only observe average outcomes for this 4 year period. Most of these years fall in the interruption period, and we therefore include the multi-year bin in the sample and start the window in 1947. The end of the window is chosen to allow held-up stations to launch and unfold their full effect. Based on the event study graph, it takes three to four years for the impact of a station to fully unfold. We therefore end the sample window in 1957, three years after the initial full-year operation of a typical held-up station.[20] The estimates again show significant negative effects and confirm that the effects arise in places were launches happen and not in places that were next in line to have a launch. The results of this experiment align closely with the baseline estimates. If anything, the point estimate for older workers is marginally larger with a 0.3 percent decline in employment for older workers. We don't want to over-interpret this difference since the baseline estimates are well within the confidence intervals of the experimental ones. The experiment suggests that the baseline results based on the overall rollout capture the causal impact of television.

3.2.1 Two-Way Fixed Effect Estimators

A recent and growing literature discusses identification strategies in two-way fixed effect settings (Borusyak and Jaravel, 2018; de Chaisemartin and D'Haultfœuille, 2020; Callaway and Sant'Anna, 2020). This literature proposes several fixes for alternative DiD related specifications. Unfortunately, there is no estimator that applies to our case yet. Our setting deviates from the canonical binary staggered DiD model in two ways: It features a dosage component where the intensity of treatment varies between units and, in addition, there are repeated treatment events. A recent review by de Chaisemartin and D'Haultfœuille (2023) describes this case as an avenue for future research. We instead follow recent applied work with a similar identification strategy (Fuest et al. (2015), Serrato and Zidar (2016), and Drechsler et al. (2017), and Schmidheiny and Siegloch (2019)) and perform robustness tests using leads and lags of the treatment variable. Schmidheiny and Siegloch (2019) show that a standard event-study style analysis can be implemented by estimating a distributed lag model of the following form:

[20] The interruption ends in 1952 and most held-up stations start broadcasting a year later in 1953 (with a few launches in 1952, 1954 and 1955). The first full year of broadcasting thus occurs in 1954. Alternative end dates deliver similar results.

$$\Delta E_{iagt} = \alpha_{gt} + \underbrace{\sum_{j=0}^{a} \beta_{g,j} \text{ Stations}_{a,t-j}}_{\text{Lagged Stations}} + \underbrace{\sum_{k=1}^{a} \beta_{g,k} \text{ Stations}_{a,t+k}}_{\text{Future Stations}} + \boldsymbol{\Pi X}_{iagt} + \epsilon_{iagt} \tag{2}$$

the $\beta_{g,j}$ coefficients capture the past impact of lagged stations and $\beta_{g,k}$ the impact of future stations. We use the approach developed in their paper to construct the event-study plots discussed above (Figure 3). Panels A and B of this figure show effects separately for two groups (g): individuals over the age of 50 and under the age of 50.

3.3 Work Hours

So far, the analysis focused on extensive margin responses, and we now additionally allow for changes to hours worked. The Social Security data does not contain information on work hours so we supplement our analysis with data from the CES. Recall that this data is aggregated at the MSA or state level and we therefore run the difference-in-differences analysis at this more aggregated level.

We first estimate employment effects in the CES data, and similar to our baseline results, find that TV station launches lead to negative employment effects (Panel A in Table 4). The smaller sample size of the CES, however, reduces the power of these estimates and the results are therefore not statistically significant. Because of the reduced sample size, we first show results that replace the year fixed effects with year trends and subsequently allow for more flexible time effects (cubic, state specific trends) and ultimately year effects. All the specifications show similar effects, with point estimates around a one percentage point decline in employment.

Panel B shows the change in total hours worked, the product of employment and average hours worked. Total hours also decline by about 1 percentage point. The drop in employment effect alone thus explains nearly all of the change in total hours worked, whereas average hours worked are unaffected by the launch of television stations. This result aligns with historical accounts of the labor market in the 1950's, when workers had only limited control over working hours and work hours were largely set through union agreements and there was minimal scope for part-time work. The extensive margin was thus the main plausible margin of adjustment and that is indeed what we find in the data.

3.4 Heterogeneous Effects: The Role of Retirement

We next unpack the employment changes into several channels and in particular turn to Costas' hypothesis about changes in retirement behaviour (Costa, 1998).

To evaluate the retirement hypothesis more directly, Figure 5 disaggregates the overall effects into three possible transition rates by age. We differentiate entries, exits and retirements (a subset of exits) and define retirement as a permanent exit from the labor force.[21] The results show a large and significant increase in retirement rates among older workers. Among the age group over 60 the probability of retirement increases roughly 0.3 percentage points, while reassuringly we find no discernible effect on the retirement of age groups below 50 (Figure 5). These retirement effects are also substantially larger than the effects on other labor market flows. Figure 5 shows only modest changes in other labor market flows, and these effects are dwarfed by the magnitude of retirement effects. Moreover, the observed increase in exit rates among older workers is overwhelmingly driven by rising retirement probabilities.

The results also align with reports that the 1950's were a period that transformed the perception of retirement. In earlier decades retirement happened when people could no longer work; in the middle of the century attitudes shifted and retirement became seen as a desirable third stage of life with additional time for leisure activities (Costa (1998)). Our finding supports Costas' hypothesis that the cheap availability of around the clock entertainment contributed to the transformation of US retirement patterns.

The retirmenet-specific results are also helpful in distinguishing the impact of the value of leisure from other potential channels. Appendix 8.3 presents a model of life-time labor supply and shows that responses to entertainment innovations are largest at the retirement margin. Other studies of Television show that watching Television also affected voting, childcare, print media, teenage pregnancy, consumption of large consumer items and household debt (Gentzkow (2006); Gentzkow and Shapiro (2008); Baker and George (2010); Campante and Hojman (2013); Thomas (2019); Kim (2020); La Ferrara et al. (2012); Angelucci et al. (2021); and Nieto (2020)). While some of these changes potentially could spill over to labor supply, most of such effects would go in the opposite direction to our finding and increase labor supply. Moreover, factors like teenage pregnancy, childcare and household debt are more concentrated among younger age groups and would struggle to explain the major changes in retirement patterns that we find.

4 Discussion

Television has accounted for a majority of Americans' leisure time since the middle of the twentieth century, making it the entertainment innovation most likely to influence labor supply trends. Several recent studies analyse the impact of entertainment technologies on labor supply through the lense of labor supply models. A recent review of the labor supply literature by Abraham and

[21] The long-run work histories of the longitudinal Social Security data allows us to observe whether individuals return to work later in life and we define retirement as permanent exits from the workforce.

Kearney (2020) points out that there is considerable uncertainty about the sensitivity of these results to modeling assumptions. The estimates in this paper can provide benchmarks for the relevant model parameters and help assess the plausibility of the prevelant modeling assumptions.

A first notable difference between our results and the results on videogames in Aguiar et al. (2021) is that we find mainly effects on older workers, while videogames mostly affected younger workers. Our findings highlight that entertainment technologies do not necessarily reduce the labor supply of the young. Their effect rather depends on the main user groups of a new technology: older individuals for TV, younger individuals for videogames. This is also relevant for discussion of alternative future technologies, their effect will likely vary and depend on the demographic group that uses these technologies. A further implication of this difference is that, even if interpreted through a model, our results may not carry over to different time periods and/or demographic groups if preference parameters are different between these settings. In our setting, older individuals are the most affected group and we will model the preferences of this group. It will be up to researchers to determine how much deviations from these values are plausible in alternative settings.

A further challenge with comparisons of different technologies is that technology shocks lack natural units. As a result, it is difficult to assess the "magnitude" of the shock and to draw generalizable conclusions from specific settings.[22] A large response to a new technology could be explained by either a major improvement in the value of leisure, or by particularly responsive leisure demand functions.

We address this challenge by deriving comparative statics results that are independent of the magnitude of shocks and can thus be used to make comparisons across settings. Appendix 7 provides a link from observable reduced form outcomes to structural preference parameters. It shows that the impact of entertainment innovation depends on two preference parameters: the elasticity of leisure demand (η_T) and the substitution elasticity between leisure activities (ϵ^c_{ii}). Entertainment innovations have a larger impact when the leisure elasticity η_T is larger or when workers substitute little between alternative leisure activities (a small ϵ^c_{ii}).

At first sight, it might seem reasonable to estimate the two elasticities directly in the data. Doing so requires measuring the percentage change in leisure value, and since the value of leisure time is typically unobserved, this is not feasible. Instead, we show that the ratio of these elasticities is equal to the ratio of two observable outcomes. Therefore, one can identify the preference parameters without the need to observe the value of leisure. The two relevant outcomes are the percent change of leisure time ($\Delta T/T$) and the change in activity i's share of total leisure time

[22] A typical approach is to use prices to measure technological progress. However, this relies on the ability of statistical agencies to construct hedonic price indexes that accurately quantify the leisure value of technological progress. At this point, it becomes a circular problem.

(Δz_i). The ratio of the two changes can be thought of as a sufficient statistic for the impact of entertainment innovations and identifies $\frac{\eta_T}{\epsilon^c_{ii}}$. This result is shown in Equation (7) in the Appendix and derived in the same section.

We first analyze the change in the average amount of leisure time among older workers ($\Delta T/T$). Converting the baseline estimate of an increase in retirement rates by 0.3 percent to a change in leisure time requires a few additional assumptions. We first convert this into a change in time-use. Using a typical work-week of 40 hours and a population of n, the retirement effect translates into $0.003*40*60*n = 7n$ minutes additional weekly leisure time per station launch. These are effects per station, and we need to account for multiple station to compute the effect of the entire roll-out. This requires an assumption on the point when new stations no longer have a meaningful impact. The baseline estimate are an average effect for the stations in the sample window, and the average number of stations in the sample is between six and seven. Combining the two implies the full roll-out increased leisure time by $6.5*7n = 45.5n$ minutes. This calculation only accounts for the stations launched during our sample window and assumes that additional stations had a negligible effect. This seems reasonable given we find deminshes effects of station launches. As alternative, we show an upper bound that assumes that up to 10 stations had negative employement effects. Which implies the roll-out increase leisure time by $10*7n = 70n$ minutes. In percentage terms, the leisure time increased between $\frac{45.5n}{31*60*n} = 2.4\%$ and $\frac{70n}{31*60*n} = 3.8\%$.[23]

Next, we quantify Δz_i, the change in the portion of leisure time devoted to home entertainment. Time use data in the pre-TV era is scarce and we do not know of a data that covers representative data before 1965.[24] An alternative approach uses data from a representative 1965 cross-sectional sample of individuals over 50. Home entertainment, excluding television, took up approximately 50% of leisure time and with television this share increases to 89%. Assuming a no TV counterfactual and no other changes to leisure times, this implies a denominator of $\Delta z_i = (0.89 - 0.50) = 0.49$.

These results together identify structural preference parameters than we can compare to the claibrations in the literature. Using the ratio of $\Delta T/T$ and Δz_i, we find $\frac{\eta_T}{\epsilon^c_{ii}} = \frac{0.022}{0.49} = 0.045$ or $= \frac{0.038}{0.49} = 0.078$ at the upper bound. This means that a leisure innovation that increases z_i by 10% leads to an increase in leisure time among older cohorts around 0.4% to 0.8%. These numbers represent the ratio of two preference parameters. To separate the two, we need further assumptions. Without income effects, η_T is proportional to the Frisch labor supply elasticity (η_L), with a proportionality factor of $\frac{L}{T} = \frac{31}{40} = 0.78$ (see Appendix 8.1). A large literature has estimated η_L and provides benchmark values for this parameter. For comparability, we follow Aguiar et

[23] This uses 31 leisure hours as baseline leisure time, following leisure time estimates in Aguiar and Hurst (2007).

[24] A 1930 survey of female "homemakers" provides some time-use information. However, it only covers non-working individuals and includes only a few dozen individuals over the age of 50, our group of interest.

al. (2021) and use their benchmark value of 1.1. Using these estimates implies value between $\epsilon_{ii}^c = (\eta_L \cdot \frac{L}{T}) \cdot 0.045^{-1} = (1.1 \cdot 0.78) \cdot 0.045^{-1} = 19$ and $\epsilon_{ii}^c = 1.1 \cdot 0.78) \cdot 0.078^{-1} = 11$. Our result implies that people substitute easily between alternative leisure activities. When new entertainment technologies arrive, people mostly substitute away from other leisure activities to spend time with the new technology.

Compared to the calibrations in Aguiar et al. (2021), our estimates imply that entertainment innovations have smaller effects on labor supply. Aguiar et al. (2021)'s calibration yields $\epsilon_{ii}^c = 1.62$, significantly smaller than our estimate of 11.[25] The calibration with low ϵ_{ii}^c yields bigger labor supply effects because people are unlikely to switch between alternative leisure activities and time with new technologies therefore has to come from reduced work-time. Conversely, our estimate of a high ϵ_{ii}^c suggests less crowding out of work-time and more of other leisure activities. Putting it all together, our estimates imply a labor supply effect of entertainment innovations that is roughly one-tenth the magnitude found in Aguiar et al. (2021).

The difference in results is due to the calibration choices, rather than assumptions about the utility function. Although our result is based on homothetic preferences, a case with the Aguiar et al. (2021) utility assumption is presented in Appendix 8.2 and delivers similar conclusions. For home entertainment specifically, a leisure luxury in the Aguiar et al. (2021) framework, introducing non-homogenous preferences further widens the gap between our estimates and previous calibration parameters. Our estimates again imply a more modest role of entertainment innovations in shaping labor supply.

5 Conclusion

Economists have recently taken an interest in the possibility that entertainment technology may affect work behavior, a hypothesis explored in the context of contemporary video games by Aguiar et al. (2021). This paper studies the single most consequential improvement in entertainment technology in the twentieth century, the introduction of television. All else equal, one would expect an increase in the utility derived from leisure time through superior entertainment to reduce labor supply, particularly for workers already on the margins of labor force participation to begin with. This paper tests this hypothesis and provides a framework to quantify the magnitude of the effects and link the results to the sizable literature on labor supply trends.

We find that TV led to statistically and economically significant declines in employment during the rollout of television. Two additional results lend confidence to our main findings. First, we are able to exploit a sharp freeze in broadcast licensing to run a series of placebo tests that help rule

[25] This calculation uses Aguiar et al. (2021)'s calibration of $\beta_i = 2.39$ for computer games (Table 6) and $\bar{\eta} = 1.09$ (p.41) and computes $\epsilon_{ii}^c = \beta_i \cdot \bar{\eta} - 1 = 2.39 \cdot 1.09 - 1 = 1.62$.

out spurious associations between TV access and employment patterns. We show that "ghost stations" whose applications for broadcasts were just denied by the FCC have no effects on work, suggesting that it was indeed TV broadcasts themselves, rather than correlated or confounding trends in economic conditions, that led to the increase in retirement. Second, the effects of TV are largest for retirement-age workers; we see no evidence that TV led younger workers to quit their jobs, but the availability of TV did increase retirement rates among older workers. This is consistent with the change in the nature of retirement documented in Costa (1998), whereby leaving one's career began to happen not only by necessity but also for the enjoyment of "golden years" of leisure, and with the fact that today, according to data from Aguiar and Hurst (2007), people in the U.S. aged over 65 spend on average four hours a day watching TV.

While research and discussion of trends in labor force participation continue to focus on labor demand topics like trade and technology, we offer novel evidence on the role of an under-explored supply-side question of technical change–how entertaining is time spent at home? TV improved the outside option for people on the margins of the labor force as it rolled out in the 1940's and 1950's. The proliferation of ever more compelling TV and of broader entertainment opportunities more generally speaks to the likely persistence and importance of these effects. Entertainment technology is of course far from the only consideration, but relative to the vast amount of time an average person spends with entertainment technology, their impact on labor markets is still a very much under researched area.

References

Abraham, Katharine G and Melissa S Kearney, "Explaining the Decline in the U.S. Employment-To-Population Ratio: a review of the evidence," *Journal of Economic Literature*, 2020, *58* (3).

Acemoglu, Daron, David Autor, and David Lyle, "Midcentury Women, War, and Wages: The Effect of Female Labor Supply on the Wage Structure at Midcentury," *Journal of Political Economy*, 2004, *112* (3), 497–551.

Aguiar, Mark and Erik Hurst, "Measuring Trends in Leisure: The Allocation of Time over Five Decades," *The Quarterly Journal of Economics*, 2007, *06* (2), 61.

_ , Mark Bils, Kerwin Kofi Charles, and Erik Hurst, "Leisure Luxuries and the Labor Supply of Young Men," *Journal of Political Economy*, 2021, *129* (2).

Angelucci, Charles, Julia Cagé, and Michael Sinkinson, "Media Competition and News Diets," *NBER Working Paper*, 2021.

Baker, Matthew J and Lisa M George, "The Role of Television in Household Debt : Evidence from the 1950's," *The B.E. Journal of Economic Analysis and Policy*, 2010, *10* (1).

Becker, Gary S., "A Theory of the Allocation of Time," *The Economic Journal*, 1965, *75* (299), 493.

Blundell, R., E. French, and G. Tetlow, *Retirement Incentives and Labor Supply*, 1 ed., Vol. 1, Elsevier B.V., 2016.

Boppart, Timo and Per Krusell, "Labor supply in the past, present, and future: A balanced-growth perspective," *Journal of Political Economy*, 2020, *128* (1), 118–157.

Borusyak, Kirill and Xavier Jaravel, "Revisiting Event Study Designs," *SSRN Electronic Journal*, 2018, pp. 1–25.

Bose, Gautam, Tarun Jain, and Sarah Walker, "Women's Labor Force Participation and Household Technology Adoption," *SSRN Electronic Journal*, 2020.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with multiple time periods," *Journal of Econometrics*, 2020, (xxxx), 1–31.

Campante, Filipe R. and Daniel A. Hojman, "Media and polarization. Evidence from the introduction of broadcast TV in the United States," *Journal of Public Economics*, 2013, *100*, 79–92.

Chadi, Adrian and Manuel Hoffmann, "Television, Health, and Happiness: A Natural Experiment in West Germany," *SSRN Electronic Journal*, 2021, (14721).

Coen-Pirani, Daniele, Alexis León, and Steven Lugauer, "The effect of household appliances on female labor force participation: Evidence from microdata," *Labour Economics*, 2010, *17* (3), 503–513.

Costa, Dora L., *The Evolution of Retirement: An American Economic History, 1880-1990*, University of Chicago Press, 1998.

De Cavalcanti, Tiago V.V. and José Tavares, "Assessing the "engines of liberation": Home appliances and female labor force participation," *Review of Economics and Statistics*, 2008, *90* (1), 81–88.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

_ **and** _ , "Two-way fixed effects and differences-in-differences with heterogeneous treatment effects : a survey," *Economic Journal*, 2023, *26.*

Della Vigna, Stefano, Ruben Enikolopov, Vera Mironova, Maria Petrova, and Ekaterina Zhuravskaya, "Cross-border media and nationalism: Evidence from serbian radio in croatia," *American Economic Journal: Applied Economics*, 2014, *6* (3), 103–132.

Drechsler, Itamar, Alexi Savov, and Philipp Schnabl, "The deposits channel of monetary policy," *Quarterly Journal of Economics*, 2017, *132* (4), 1819–1876.

Durante, Ruben, Paolo Pinotti, and Andrea Tesei, "The political legacy of entertainment TV," *American Economic Review*, 2019, *109* (7), 2497–2530.

Enikolopov, Ruben, Maria Petrova, and Ekaterina Zhuravskaya, "Media and political persuasion: Evidence from Russia," *American Economic Review*, 2011, *101* (7), 3253–3285.

Fisher, Franklin M, "Community Antenna Television Systems and the Regulation of Television Broadcasting," *American Economic Review*, 1966, *56* (2), 320–330.

Fisher, Kimberly, Jonathan Gershuny, Sarah M. Flood, Joan Garcia Roman, and Sandra L. Hofferth, "American Heritage Time Use Study Extract Builder: Version 1.2," *IPUMS*, 2018.

French, Eric, "The effects of health, wealth, and wages on labour supply and retirement behaviour," *Review of Economic Studies*, 2005, *72* (2), 395–427.

_ **and Bailey J Jones**, "The Effects of Health Insurance and Self-Insurance on Retirement Behavior," *Econometrica*, 2011, *79* (3), 693–732.

Fuest, Clemens, Andreas Peichl, and Sebastian Siegloch, "Do Higher Corporate Taxes Reduce Wages?," 2015, (9606).

Gelber, Alexander, "Taxation and The Earnings of Husbands and Wives: Evidence from Sweden," *The Review of Economics and Statistics*, 2014, *96* (2), 287–305.

Gennaro, Gloria and Elliott Ash, "Emotion and Reason in Political Language," *Economic Journal*, 2022, *132* (643), 1037–1059.

Gentzkow, Matthew, "Television and voter turnout," *The Quarterly Journal of Economics*, 2006, *121* (August), 931–971.

— and Jesse Shapiro, "Preschool Television Viewing and Adolescent Test Scores: Historical Evidence from the Coleman Study," *Quarterly Journal of Economics*, 2008, *123* (February), 279–323.

González-Chapela, Jorge, "On the price of recreation goods as a determinant of male labor supply," *Journal of Labor Economics*, 2007, *25* (4), 795–824.

Greenwood, Jeremy, Ananth Seshadri, and Mehmet Yorukoglu, "Engines of liberation," *Review of Economic Studies*, 2005, *72* (1), 109–133.

— **and Guillaume Vandenbroucke**, "Hours Worked (Long-Run Trends)," *The New Palgrave Dictionary of Economics*, 2008, (September), 75–81.

—**, Nezih Guner, Georgi Kocharkov, and Cezar Santos**, "Technology and the changing family: A unified model of marriage, divorce, educational attainment, and married female labor-force participation," *American Economic Journal: Macroeconomics*, 2016, *8* (1), 1–41.

Haines, Michael R., "Historical, Demographic, Economic, and Social Data: The United States, 1790-2002," Technical Report, Inter-university Consortium for Political and Social Research 2010.

Jensen, Robert and Emily Oster, "The power of TV: Cable television and women's status in India," *Quarterly Journal of Economics*, 2009, *124* (3), 1057–1094.

Juhn, Chinhui and Simon Potter, "Changes in labor force participation in the United States," *Journal of Economic Perspectives*, 2006, *20* (3), 27–46.

Kearney, Melissa S. and Phillip B. Levine, "Media influences on social outcomes: The impact of MTV's 16 and pregnant on teen childbearing," *American Economic Review*, 2015, *105* (12), 3597–3632.

— **and Phillip B Levine**, "Early Childhood Education by MOOC: Lessons from Sesame Street," *American Economic Journal: Applied Economics*, 2019, *11* (1), 318–350.

Kim, Woojin, "Television and American Consumerism," *Working Paper*, 2020.

Koenig, Felix, "Technical Change and Superstar Effects: Evidence from the Rollout of Television," *American Economic Review: Insights*, 2023.

Kopecy, Karen, "The Trend in Retirement," *International Economic Review*, 2011, *52* (2), 287–316.

Kopytov, Alexandr, Nikolai Roussanov, and Mathieu Taschereau-Dumouchel, "Cheap Thrills: The Price of Leisure and The Global Decline in Work Hours," *American Economic Journal: Macroeconomics*, 2023.

La Ferrara, Eliana, Alberto Chong, and Suzanne Duryea, "Soap operas and fertility: Evidence from Brazil," *American Economic Journal: Applied Economics*, 2012, *4* (4), 1–31.

Lazear, Edward P., "Retirement from the labor force," *Handbook of Labor Economics*, 1986, *1* (C), 305–355.

Long, Wayne, "Employment Covered Under the Social Security Program," *Social Security Bulletin*, 1988, *51* (10).

Lumsdaine, Robin L. and Olivia S. Mitchell, "New developments in the economic analysis of retirement," *Handbook of Labor Economics*, 1999, *3 PART* (3), 3261–3307.

Mastrorocco, Nicola and Luigi Minale, "News media and crime perceptions: Evidence from a natural experiment," *Journal of Public Economics*, 2018, *165*, 230–255.

Mcgrattan, Ellen R and Richard Rogerson, "Changes in Hours Worked, 1950-2000," *Federal Reserve Bank of Minneapolis Quarterly Review*, 2004, *28* (1), 14–33.

Ngai, L. Rachel and Barbara Petrongolo, "Gender Gaps and the rise of the service economy," *American Economic Journal: Macroeconomics*, 2017, *9* (4), 1–44.

Nieto, Adrián, "Television, time use and academic achievement: Evidence from a natural experiment," *working paper*, 2019, (19).

Nieto, Adrian, "Television and the Labour Supply: Evidence from the Digital Television Transition in the UK," *working paper*, 2020, (April).

of the Census, United States Bureau, "Current Population Survey (CPS)/Social Security Summary Earnings (SER) Exact Match File, March 1978," Technical Report, Inter-university Consortium for Political and Social Research 2009.

_ , "County and City Data Book [United States] Consolidated File: County Data, 1947-1977," Technical Report, Inter-university Consortium for Political and Social Research 2012.

Olken, Benjamin A., "Do television and radio destroy social capital? Evidence from indonesian villages," *American Economic Journal: Applied Economics*, 2009, *1* (4), 1–33.

Putnam, Robert, "Bowling Alone: America's Declining Social Capital," *Journal of Democracy*, 1995, *6* (1), 65–78.

Rachel, Łukasz, "Leisure-Enhancing Technological Change," *working paper*, 2020.

Schmidheiny, Kurt and Sebastian Siegloch, "On Event Study Designs and Distributed-Lag Models: Equivalence, Generalization and Practical Implications," *CESifo Working Papers*, 2019, *7481* (12079).

Serrato, Juan Carlos Suárez and Owen Zidar, "Who benefits from state corporate tax cuts? A local labor markets approach with heterogeneous firms," *American Economic Review*, 2016, *106* (9), 2582–2624.

Television-Factbook, "Vol 5, 9, 11, 13, 15, 17, 19, 21, 23, 25, 27, 29, 31," Technical Report.

Thomas, Michael, "Was Television Responsible for a New Generation of Smokers?," *Journal of Consumer Research*, 2019, *46* (4), 689–707.

Vandenbroucke, Guillaume, "Trends in hours: The U.S. from 1900 to 1950," *Journal of Economic Dynamics and Control*, 2009, *33* (1), 237–249.

Wang, Tianyi, "Media, Pulpit, and Populist Persuasion: Evidence from Father Coughlin," *SSRN Electronic Journal*, 2020.

Wasserman, Max J and John R Arnold, "Old-Age Insurance: Covered Workers and Average and Median Taxable Wages in 1937," *Social Security Bulletin*, 1939, *2* (6).

Yanagizawa-Drott, David, "Propaganda and Conflict: Theory and Evidence From the Rwandan Genocide," *Quarterly Journal of Economics*, 2014, *129* (4), 1947–1994.

6 Figures and Tables

Figure 1: Hours of Television Watching per Day

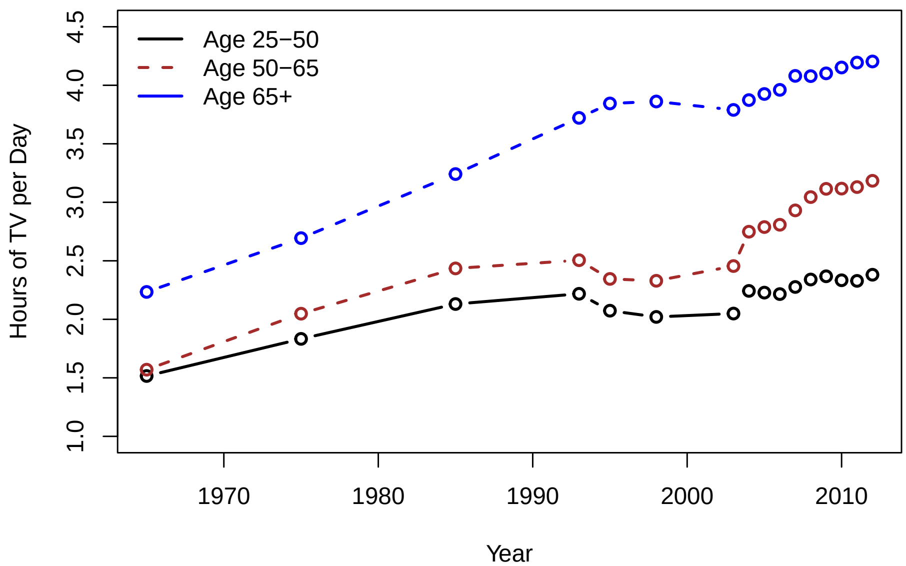

Notes: The figure shows the amount of time American's spend watching television as primary activity. Data are rolling averages from the Historic American Time Use Study (AHTUS). The hours refer to "primary activity."

Figure 2: ITM-Measured Signal Strength in 1950

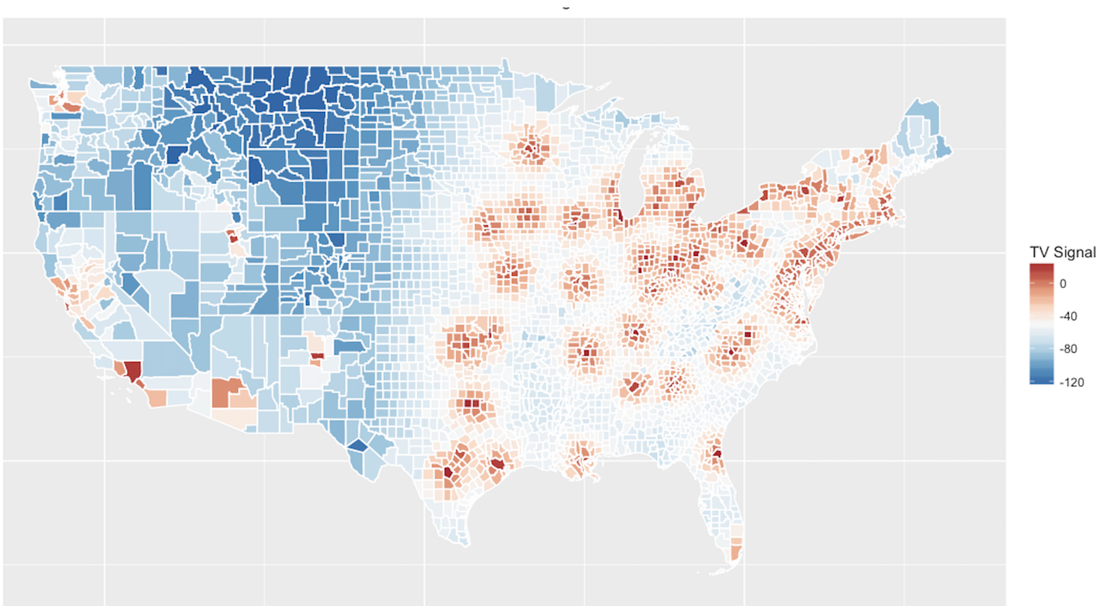

Notes: The figure shows the signal level, in decibels, of the strongest station in each county in 1950, as computed with the ITM. Broadly, counties shaded red had TV access, while counties shaded blue did not; signals whose strength was less than -50 decibels, where the map turns from red to blue, were effectively unwatchable. Not shown in this visualization of the data is the *number* of stations available locally.

Figure 3: Effect of Television Launches by Age Groups

(a) Individuals Aged over 50

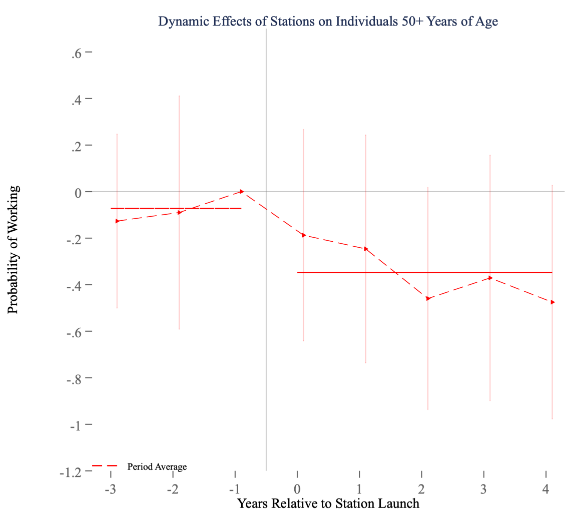

(b) Individuals Aged under 50

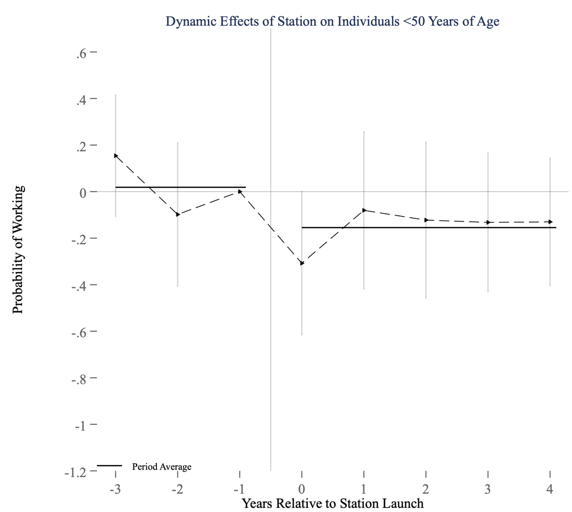

Notes: The figure shows dynamic effects of the launch of a TV station, separately for Individuals over and under 50 years of age. Specifically, these are based on individual level regression of a dynamic version of equation 1 (adding leads and lags of the treatment variable) and the approach in Schmidheiny and Siegloch (2019) and related recent work with a continuous treatment variable. The figure shows 95 percent confidence intervals. See the text for further details.

Figure 4: Balance Test for TV Treatment

(a) DiD Sample

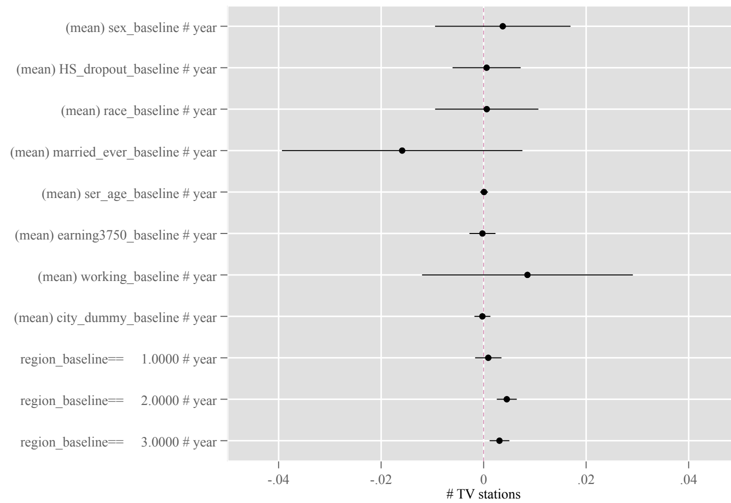

(b) Interruption Experiment Sample

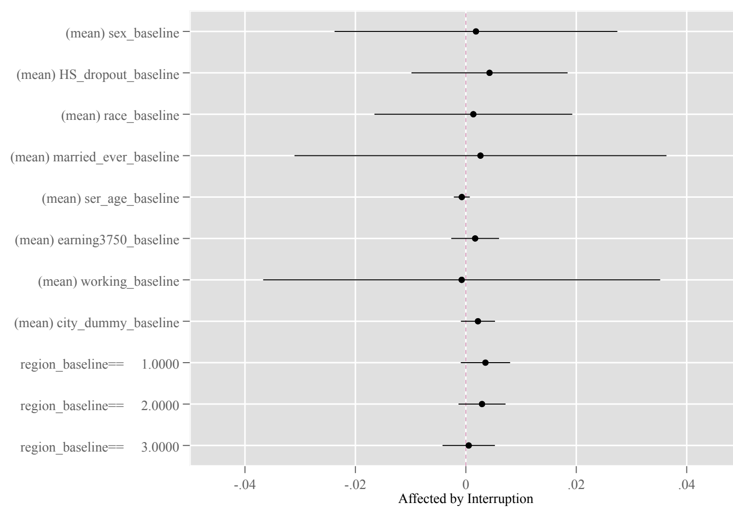

Notes: Panel A shows coefficients from a local area level regression, regressing TV stations on linear trends in baseline demographics and year and location fixed effects. Baseline demographics are based on the earliest period in the data (the multi-year bin 1937-1946). Panel B analyses the locations in the "interruption experiment" sample. The results are based on a regression of an indicator for being affected by the interruption on local baseline demographics. Standard errors are clustered at the local area level, observations are weighted by cell-size and 99 percent confidence intervals are shown. We convert units of continuous variable (average earnings and average age) into standard deviations to make the magnitudes of coefficients comparable. See the text for further details.

Figure 5: Effects of TV on Entry, Exit, and Retirement

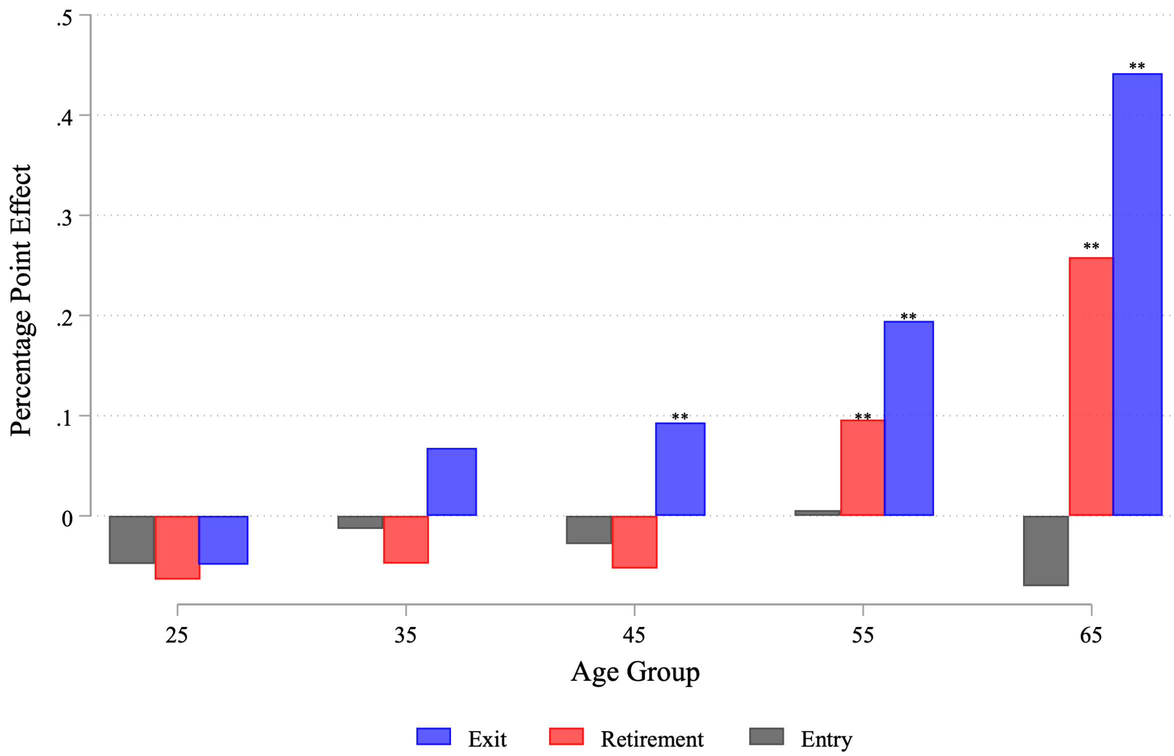

Notes: The figure shows the impact of television on job transitions. Effects on employment entry are shown in black, on exits in blue and on retirement in red. Retirement is an indicator with value hundred in the year a worker exits permanently from employment (proxied by the absence of a work observation until the end of our data). The plotted results are coefficients from the baseline difference-in-difference regressions from Table 1 run separately for the three outcome variables (exit, entry, retirement) and allowing for separate coefficients by age group. The regression uses data between 1951-1960 when data is annual and annual flows can be calculated and is based on 293,431 observations. The x-axis shows the mid point of ten year age bins, e.g., 55 represents ages 50 to 59. For additional specification details see Table 1, column 2. $**$ indicates that the coefficient is significant at the 5 percent level.

Table 1: Individual-level Effects of TV on Employment

	(1)	(2)	(3)	(4)	(5)
	$\Delta Employment$				
Effect of TV stations $\times$					
Age over 50	-0.276***	-0.241***	-0.235***	-0.234***	-0.243***
	(0.0683)	(0.0698)	(0.0740)	(0.0758)	(0.0751)
Age under 50	-0.0911	-0.101*	-0.0773	-0.0773	-0.0853
	(0.0590)	(0.0589)	(0.0630)	(0.0631)	(0.0638)
Female	0.0250	0.0251	0.0233	-0.0181	0.0241
	(0.0831)	(0.0829)	(0.0829)	(0.0946)	(0.0830)
Observations	292,448	292,448	292,448	292,448	292,448
R-squared	0.048	0.048	0.048	0.048	0.048
Year FE	Yes	Yes	Yes	Region X Year	Yes
Person FE	Yes	Yes	Yes	Yes	Yes
Age FE	No	Yes	Yes	Yes	Yes
Quartic Trends	No	No	Region	No	Demographics
Mean DV age over 50	0.337	0.337	0.337	0.337	0.337
Mean DV age under 50	1.421	1.421	1.421	1.421	1.421

Notes: The table shows individual level regressions of $\Delta Employment$ with value 100 for a labor market entry and -100 for a labor market exit on the number of TV stations available in the local area. The effect of stations is interacted with indicators for Age over 50 (specifically, 50+), Age under 50, and Female. Data are at the individual level and covers individuals over the age of 21 and spans 1937-1960, at annual frequency from 1951 onward and multi-year averages for earlier periods (see text for details). All regressions include gender-specific year fixed effects. Trends use a 4th order polynominal in time. Regions are census regions. Demographics are indicators for ever married, census regions, and local employment rates in the baseline period. Television is measured at the MSA level. Standard errors are clustered at the same level and span 134 clusters. Observations are weighted using the strategy described in the text. Source: SSA-CPS employment records and Television Factbooks $***p < 0.01, **p < 0.05, *p < 0.1$

Table 2: Diminishing Effects of additional TV stations

	(1)	(2)	(3)	(4)	(5)
	Diminishing Effect		Using only first Station(s)		
	quadratic	logarithmic	[0-1]	[0-5]	[all]
	Panel A: Age over 50				
TV stations	-0.652***		-0.919	-0.602***	-0.261***
	(0.113)		(1.056)	(0.115)	(0.0626)
TV stations^2/10	0.218***				
	(0.0574)				
log(TV stations + 1)		-2.211***			
		(0.348)			
	Panel B: Age under 50				
TV stations	-0.227***		-0.364	-0.0865	-0.0765
	(0.0838)		(0.512)	(0.0785)	(0.0553)
TV stations^2/10	0.0773**				
	(0.0335)				
log(TV stations + 1)		-0.840***			
		(0.279)			
Observations	292,448	292,448	292,448	292,448	292,448
R-squared	0.048	0.048	0.048	0.048	0.048

Notes: The table shows how the effect of new TV station diminishes when they launch in an area where stations are already active. Column 1 allows for effects to decline quadratically and 2 uses $log(TV + 1)$ as regressor, and thus assumes effects decline logarithmically. Columns 3 through 5 estimate the effect of TV using only the first launches in an area, with the number of launches used indicated in the column title. The estimates use the baseline specification in column 2 of Table 1. $***p < 0.01, **p < 0.05, *p < 0.1$

Table 3: Effects of TV on Employment Using Variation from Regulator Shutdown

	(1)	(2)	(3)	(4)	(5)	(6)
			Interruption Experiments			
	Placebo Test		Pre vs. Post		Pre vs. Next Ranked	
			Panel A: Age over 50			
TV Stations	-0.240***	-0.219***	-0.169***	-0.191***	-0.283***	-0.290***
	(0.0668)	(0.0711)	(0.0609)	(0.0669)	(0.0949)	(0.103)
Blocked stations	-0.0019	0.0488				
	(0.0650)	(0.0687)				
			Panel B: Age under 50			
TV Stations	-0.118**	-0.082	0.009	-0.006	-0.091	-0.096
	(0.058)	(0.062)	(0.054)	(0.062)	(0.076)	(0.085)
Blocked stations	-0.062*	-0.004				
	(0.036)	(0.040)				
Observations	285,147	285,147	231,931	231,931	151,691	151,691
R-squared	0.048	0.048	0.047	0.047	0.071	0.071
Region Trends	No	Yes	No	Yes	No	Yes

Notes: The table shows the impact of television on employment rates, using variation from the regulator shut-down. Columns 1 and 2 compare the effect of TV stations and stations that were blocked during the regulator shutdown 1948-1952. Columns 3 through 6 focus on variation from the rollout interruption. Experiment "Pre vs. Post" in Column 3 and 4 uses the variation pioneered by Gentzkow (2006) and compares locations with TV station launches before and after the interruption start and end date (1947-1954). Columns 5 and 6 ("Pre vs. Next Ranked") uses locations ranked next in the FCC priority ranking as the control group and focus on sample years when TV timing was driven by the interruption (years of the interruption and unwind, 1947-1957). The estimates use the baseline specification in column 2 of Table 1. See Table notes for additional details $***p < 0.01, **p < 0.05, *p < 0.1$

Table 4: MSA-Level Effects of TV on Employment and Hours in Manufacturing

	(1)	(2)	(3)	(4)
	Panel A: CES Log Manufacturing Employment			
Stations	-0.011	-0.012	-0.012	-0.011
	(0.0085)	(0.0084)	(0.0085)	(0.011)
Observations	446	446	446	446
R-squared	0.994	0.994	0.997	0.994
	Panel B: CES Log Total Manufacturing Hours			
Stations	-0.012	-0.013	-0.013	-0.010
	(0.0084)	(0.0083)	(0.0087)	(0.011)
Observations	446	446	446	446
R-squared	0.993	0.993	0.997	0.994
Area Effects	Yes	Yes	Yes	Yes
Trends	Yes	Cubic	State	No
Year Effects	No	No	No	Yes

Notes: The table shows regressions of labor market outcomes on the number of TV stations available. Data are at the MSA level. Specifically, the outcomes are log employment and log total hours from the CES manufacturing data, respectively. Robust standard errors in parentheses. $***p < 0.01, **p < 0.05, *p < 0.1$

ONLINE APPENDIX

7 Appendix C: Representative Agent Framework

We provide a representative agent framework that closely mirrors prior work in the area and builds on the model Aguiar et al. (2021) use to study the effect of videogames.[26] The model helps understand how aggregate labor supply in the economy (i.e. total work hours) responds to innovations that enhance the value of leisure time.

Take a representative agent who derives utility from consumption c and a variety of leisure activities α_i:

$$u(c, v(\alpha_1, ..., \alpha_n)) \tag{3}$$

where $v(\alpha_1, ..., \alpha_n)$ is homogenous of degree one and thus nests a wide range of standard utility functions, e.g. Cobb-Douglas and CES utilities. Aguiar et al. (2021) use a utility function that is not homogenous of degree one, and we explore the impact of such an approach in Appendix 8.2, finding small effects on the predictions we study. We therefore focus on the simpler homogenous utility case first.

The utility maximisation problem of this agent is:

$$\max_{a_i, c} \quad u(c, v(\alpha_1, ..., \alpha_n)) \quad \text{s.t. } L + T \leq 1, \sum_{i=1}^{n} t_i \leq T, c \leq w \cdot L, \alpha_i = \theta_i \cdot t_i \tag{4}$$

Time can be spend either working (L) at wage rate w or on leisure (T) and we normalise total time to 1 (for simplicity, we do not separately model home-production). Leisure time is split between the i different types of activities ($\sum_i t_i = T$). To introduce a role for technological progress, assume that a time investment of t_i yields θ_i quality units of α_i: $\alpha_i = \theta_i \cdot t_i$. Improvements in technology increase the value of θ_i and thus the value of spending time on this leisure activity.[27] The derivations will mainly focus on the "demand for leisure" but notice that the elasticity of leisure demand is proportional to the labor supply elasticity and the discussion could thus equivalently be framed in terms of labor supply elasticities, so we use the two terms interchangeably.[28] In addition, it will be useful to define the compensated demand elasticity for activity i: $\frac{\partial ln(\alpha_i)}{\partial ln(\tilde{\theta}_i)} \equiv \epsilon_{ii}^c - 1$, where ϵ_{ii}^c is a constant for CES, Cobb-Douglas and other utility functions with a constant demand elasticity and will otherwise vary. This elasticity captures how improvements in the value of leisure from activity i affect the demand for activity i.

Solving the utility optimisation, we can derive the following two results for the impact of

[26] Similarly, a large literature in macro-economics uses this approach (for recent examples on labor supply, see Rachel (2020); Boppart and Krusell (2020))

[27] Note that the impact of technology in this model is isomorphic to a shift in preferences towards activity i. In line with most economic work, we assume that preferences are fixed throughout time.

[28] From $L + T = 1$ it follows that the elasticity of leisure demand (η_T) is proportional to the labor supply elasticity (η_L) with proportionality factor $\frac{-T}{L}$: $\eta_L = \eta_T \frac{-T}{L}$.

technical change θ_i. The impact of θ_i on T and on the time share of activity i ($z_i \equiv \frac{t_i}{T}$) are respectively (see Appendix 7.1 for derivations):

$$\frac{\partial T/T}{\partial \theta_i} = -\frac{z_i}{\theta_i}\eta_T \tag{5}$$

$$\frac{\partial z_i}{\partial \theta_i} = -\frac{z_i}{\theta_i}\epsilon_{ii}^c \tag{6}$$

Equation (5) shows the impact of leisure innovation on total leisure time, and η_T denotes the elasticity of demand for leisure (T). The impact of an entertainment innovation on T depends on the activity's productivity (θ_i), the time share of the affected leisure activity (z_i) and the elasticity of leisure demand (η_T). A challenge with this result is that the effects depend on the technology (θ_i), and technology has no natural unit. This makes estimates hard to interpret without a credible cardinal measure for θ_i. It might look as if we could avoid this issue by bringing θ_i to the left side of the equation. However, this implies that we can no longer estimate the left hand side. The LHS would be the percent change in leisure time in response to a given percent change in the value of leisure. Since we don't know the latter, we cannot estimate this object.

We propose an alternative approach that is independent of technology units and expresses individual responses as "time-use elasticities." These elasticities normalise the reduced form change in work time by a numeraire that captures the magnitude of the technology shock and in our case the numeraire is z_i – the time share of leisure activities i. z_i provides a revealed preference metric for the quality improvement that results from a new technology since people will spend more time with a more appealing entertainment activity. The impact of θ_i on z_i is shown in equation 6 and depends on θ_i, z_i and ϵ_{ii}^c. This impact thus depends on the unobserved units of the technology (θ_i) which makes it hard to link changes in z_i to preference parameters. Combining 5 and 6, we can cancel (θ_i) and get an expression that links preference parameters to measurable quantities only:

$$\frac{\frac{\partial T/T}{\partial \theta_i}}{\frac{\partial z_i}{\partial \theta_i}} = \frac{\partial T/T}{\partial z_i} = \eta_T/\epsilon_{ii}^c. \tag{7}$$

The relative change of T and z_i is independent of θ_i and depends only on the substitution elasticity between leisure activities (ϵ_{ii}^c), and the elasticity of leisure demand (η_T). The impact of an entertainment innovation on labor supply thus depends the relative magnitude of the two preference parameters ϵ_{ii}^c and η_T.[29] The expression holds independent of technology specific parameters and once η_T/ϵ_{ii}^c is estimated one can use the result to quantify the impact of entertainment innovations in other contexts.

[29]Quantifying the impact of other entertainment innovation is straight forward for preferences with a constant ϵ_{ii}^c. With varying ϵ_{ii}^c it requires additional information on how much ϵ_{ii}^c changes across contexts. In our empirical application, we find ϵ_{ii}^c to be fairly stable across settings and thus suggestive evidence that ϵ_{ii}^c is constant.

An appealing feature of equation (7) is that the left hand side can be estimated with simple reduced form techniques. It requires an estimate of the percent change of leisure time ($\Delta T/T$) and the change of activity i's leisure time share (Δz_i). The ratio of the two estimates can be thought of as a sufficient statistic for the impact of entertainment innovations. Deriving this result required relatively few assumptions on the structural form of the utility function. It applies to all utility functions of the form $u(c, v)$, where v is a homogenous function. Also note that equation 7 is i specific, which means that the impact of new technology may be different for innovations in different leisure categories. New technology in home entertainment could have a different effect than a similar innovation in sport activities, and there is a separate sufficient statistic for each group of leisure activities.

7.1 Derivation

We can solve the utility maximisation problem in 4 in two-steps.[30] The first step determines how a given amount of leisure time $T = \sum_1^n t_i$ is split across all leisure activities t_i by solving the lower nest optimisation problem:

$$V(T, \boldsymbol{\theta}) = \max_{\alpha_i} \quad v(\alpha_1, ..., \alpha_n) \quad \text{s.t. } T = \sum_1^n \alpha_i/\theta_i \tag{8}$$

Note that we take T as given in this step and assumed that all constraints are binding to obtain $T = \sum_1^n \alpha_i/\theta_i$. The set-up is now a standard demand derivation for good α_i, with price $\tilde{\theta}_i \equiv \theta_i^- 1$, so we can express demand for α_i in terms of $\tilde{\boldsymbol{\theta}}$.

The second step of the utility maximisation problem maximises the upper nest of the utility function. Using the indirect utility function $V(T, P) = v(\alpha_1^*, ..., \alpha_i^*)$ and defining $P = f(\boldsymbol{\theta})$ as the price index for a unit of V, the budget constraint for the upper level problem is $\frac{c}{w} + P \cdot V = 1$.[31] It can be shown that P is independent of the choice variable V when utility is homothetic.

Moreover, with $v(\boldsymbol{\alpha})$ homethetic, we can divide the indirect utility function ($V(T, P)$) by total leisure time T and write:

$$V(1, P(\tilde{\theta}_1/T, ..., \tilde{\theta}_n/T)) \tag{9}$$

where we use that we can represent the same preferences with any monotonic transformation of $v(\boldsymbol{\alpha})$. For instance, if the utility function is homogenous of degree β, we can represent the same preferences by either $V(1, P(\tilde{\theta}_1/T, ..., \tilde{\theta}_n/T))(\frac{1}{T})^\beta$ or $V(1, P(\tilde{\theta}_1/T, ..., \tilde{\theta}_n/T))$.[32]

[30] The proof of this claim is available upon request.

[31] This constraint combines the three constraints: $c \leq w \cdot L$, $L + T \leq 1$ and $P \cdot V \leq T$ and uses that they bind with equality.

[32] We can use a similar monotonic transformation to obtain a homogenous function representation of the preferences

Next we use Roy's identity to derive the demand for α_i. To do so, we first derive the following two results from the indirect utility function and differentiate 9: $\frac{\partial V}{\partial \tilde{\theta}_i} = \frac{1}{T}\frac{\partial V}{\partial P}\frac{\partial P}{\partial \tilde{\theta}_i}$ and $\frac{\partial V}{\partial T} = -\frac{\kappa}{T^2}$ where $\kappa = \sum_{i=1}^{n} \frac{\partial V}{\partial P}\frac{\partial P}{\partial \tilde{\theta}_i}\tilde{\theta}_i$. Using these results in Roy's identity, we get:

$$\alpha_i = -\frac{\frac{\partial V}{\partial \tilde{\theta}_i}}{\frac{\partial V}{\partial T}} = \frac{T \cdot \partial V/\partial P \cdot \partial P/\partial \tilde{\theta}_i}{\kappa} = \Pi_i T \tag{10}$$

where the last step combines all the terms that are independent of T into $\Pi_i \equiv \frac{\partial P/\partial \tilde{\theta}_i}{\sum_{i=1}^{n} \frac{\partial P}{\partial \tilde{\theta}_i}\tilde{\theta}_i} = \frac{\partial ln(P)}{\partial \tilde{\theta}_i}$, where the last equality follows from the fact that P is homogenous of degree 1 and hence $P = \sum_{i=1}^{n} \frac{\partial P}{\partial \tilde{\theta}_i}\tilde{\theta}_i$. Recall that P is the utility value of a unit of leisure time, and Π_i thus measures how entertainment innovation affect P.

Finally, we derive an expression for $z_i \equiv t_i/T$, the share of leisure time spend on activity i . Cross-multiplying 10 with $\frac{\tilde{\theta}_i}{T}$ and using $t_i = \alpha_i \cdot \tilde{\theta}$, we find z_i:

$$z_i = \frac{\alpha_i \cdot \tilde{\theta}}{T} = \Pi_i \tilde{\theta}_i = \frac{\partial P}{\partial \tilde{\theta}_i}\frac{\tilde{\theta}_i}{P} \tag{11}$$

where the final equality uses the definition of Π_i and the fact that P is homogenous of degree 1. The time share of activity i is thus equal to P's price elasticity and hence independent of T. The level of available leisure time thus does not impact the distribution of leisure time across activities. In our homothetic preference setting, leisure shares only change with technology and z_i therefore provides a revealed preference metric for the degree of technical change in the economy. The general insight holds more broadly, however, with non-homothetic preferences additional modification is necessary. If preferences are non-homothetic z_i expands or falls with available leisure time depending on whether activity i is a leisure "luxury" or "necessity." One can infer technical change from changes in z_i after controlling for the "Engel's curve" expansion path of the activity (see Aguiar et al. (2021) for a related approach).

8 Demand Response to Technical Change

We now turn to a technical change that affects $\tilde{\theta}_i$. First, consider how $\tilde{\theta}_j$ affects the total amount of leisure time T:

$$\frac{\partial T}{\partial \tilde{\theta}_j}\frac{\tilde{\theta}_j}{T} = \frac{\partial T}{\partial P}\frac{P}{T}\frac{\partial P}{\partial \tilde{\theta}_j}\frac{\tilde{\theta}_j}{P} = z_j \eta_T \tag{12}$$

where the last equality uses 11 and defines the elasticity of leisure to P as $\eta_T \equiv \frac{\partial T}{\partial P}\frac{P}{T}$. Using $\frac{\partial T}{\partial \tilde{\theta}_j} = -\frac{\partial T}{\partial \theta_j}\theta_j^2$ and multiplying both sides with $\frac{T}{\theta_j}$ yields 5 in the main text.

represented by the homothetic function $v(\boldsymbol{a})$.

A practical challenge with estimating $\frac{\partial T}{\partial \tilde{\theta}_j}\frac{\tilde{\theta}_j}{T}$ is that it requires information about the change in θ, which is typically unobserved. For instance, when technical change drives changes in $\tilde{\theta}_j$, we typically do not know the resulting percent change in $\tilde{\theta}_j$ and we can thus only estimate the reduced form expression: $\sigma_{T_j} = \eta_{T_j}\tilde{\theta}_j$, which depends on the magnitude of the shock ($\tilde{\theta}_j$) and thus difficult to compare across settings.

To make further progress, we need a scalar that captures the magnitude of technical change. To do so, consider the demand response *within* the inner nest. That is, the demand for each of the individual leisure activities. We will focus on the compensated elasticity and hold total time T constant in the derivation. The definition of $\alpha_i = t_i/\tilde{\theta}_i$ implies $\frac{\partial ln(t_i)}{\partial ln(\tilde{\theta}_i)} = 1 + \frac{\partial ln(\alpha_i)}{\partial ln(\tilde{\theta}_i)} \equiv \epsilon^c_{ii}$ and denote this compensated time use elasticity for acitivty i by ϵ^c_{ii}. Combining this result with the derivative of the log of 10 yields a solution for ϵ^c_{ii}:

$$\epsilon^c_{ii} = 1 + \frac{\partial \Pi_i}{\partial \tilde{\theta}_i}\frac{\tilde{\theta}_i}{\Pi_i} \tag{13}$$

Π_i embodies the marginal value of leisure and the result shows that the elasticity of Π_i is equal to $\epsilon^c_{ii} - 1$. Entertainment innovation therefore lead to a larger change in the value of leisure if demand for activities is elastic. Conversely, if time spending is inelastic, people shift less time towards the improved activity and a shift in technology has a smaller impact on the value of leisure. The impact of entertainment innovations thus depends on the demand elasticity for leisure activities.

Next, consider the effect of $\tilde{\theta}_i$ on z_i. Differentiating 11 and using 13 and 11 we find:

$$\frac{\partial z_i}{\partial \tilde{\theta}_i} = \Pi_i + \frac{\partial \Pi_i}{\partial \tilde{\theta}_i}\tilde{\theta}_i = \Pi_i \epsilon^c_{ii} = z_i/\tilde{\theta}_i \epsilon^c_{ii} \tag{14}$$

This expression shows the effect of technical change on z_i. Since compensated demand is downward sloping ($\epsilon^c_{ii} < 0$) and hence an entertainment innovation ($\tilde{\theta} \downarrow$) increases z_i. The magnitude of the change depends on the substitution elasticity (ϵ^c_{ii}) and the size of the technology change.

We can now derive equation 6 in the main text by multiplying both sides of 14 with $\frac{1}{\theta_j^2}$ and using the fact that $\frac{\partial z_i}{\partial \tilde{\theta}_j} = -\frac{\partial z_i}{\partial \theta_j}\theta_j^2$.

8.1 The Elasticity of Leisure and Labor

If we are further interested in separately identifying η_T and ϵ^c_{ii}, we need additional information on the utility function. The time constraint $T = 1 - L$ and $L = c/w$ implies $T = 1 - c/w$ and hence

$$\eta_T = \frac{\partial T}{\partial P}\frac{P}{T} = -\frac{\partial c}{\partial P}\frac{1}{w}\frac{P}{T}. \tag{15}$$

The Slutsky equation implies:

$$\frac{\partial c}{\partial P} = \frac{\partial c}{\partial P}|_C + c\frac{\partial c}{\partial m} = \frac{\partial V}{\partial 1/w}|_C + c\frac{\partial c}{\partial m}. \tag{16}$$

where $\frac{\partial c}{\partial m}$ is the income effect and $|_C$ indicates the compensated demand change. The last equality uses the symmetry of the substitution matrix. Furthermore, differentiating the total time constraint $L + P \cdot V = 1$ and using that P is independent of w ($\frac{\partial P}{\partial w} = 0$) , we can show that:

$$\frac{\partial V}{\partial 1/w} = -\frac{w^2}{P}\frac{\partial L}{\partial w} \tag{17}$$

note additionally that if the constraint $c = wL$ binds, then $\frac{\partial c}{\partial m} = w\frac{\partial L}{\partial m}$. And we can substantially simplify equation 16 by assuming $\frac{\partial L}{\partial m} = 0$, in line with a sizable labor supply literature that finds modest income effects.[33] Using this in 16 and combining the results with 15 and 17, we get:

$$\eta_T = -[\frac{\partial V}{\partial 1/w}|_C + c\frac{\partial c}{\partial m}]\frac{1}{w}\frac{P}{T} = \frac{\partial L}{\partial w}\frac{w}{L}\frac{L}{T} = \eta_L\frac{L}{T}$$

the elasticity of leisure time is thus the same as the compensated labor supply elasticity ($\eta_L = \frac{\partial L}{\partial w}\frac{w}{L}$) times the ratio of labor and leisure time.

8.2 Non-Homothetic Preferences

We now explore the imapct of relaxing the homotheticity assumption on $v(a_1, ..., \alpha_i)$. We follow the approach in Aguiar et al. (2021) and assuming the following functional form:

$$v(\alpha_1, ..., \alpha_i) = \sum_i \frac{\alpha_i^{1-1/\eta_i}}{1 - 1/\eta_i} \tag{18}$$

Here η_i varies across activities. The following Lagrangian determines the allocation of leisure time across all different activities:

$$\mathcal{L} = v(\alpha_1, ..., \alpha_i) - \mu\left[\sum \frac{\alpha_i}{\theta_i} - T\right] \tag{19}$$

Maximizing $\mathcal{L}$ and using $\alpha_i = t_i\theta_i$ yields the following two optimality conditions:

$$\begin{aligned} FOC_i: &\quad ln(t_i) = (\eta_i - 1) \cdot ln(\theta_i) - \eta_i \cdot ln(\mu) \\ FOC_\mu: &\quad T = \textstyle\sum_i t_i \end{aligned} \tag{20}$$

[33] note that $L + T = 1$ and $c/w + PV = 1$ imply that assumption $\frac{\partial L}{\partial m} = 0$ entails $\frac{\partial T}{\partial m} = \frac{\partial V}{\partial m} = 0$ and compensated and uncompensated demands coincide $\frac{\partial V}{\partial 1/w} = \frac{\partial V}{\partial 1/w}|_C$.

The first expression gives the demand for time t_i with activity i and the second expression is the time constraint.

Combining the two FOCs and differentiating with respect to $ln(\theta_i)$ while holding T constant yields:

$$\frac{\partial ln(\mu)}{\partial ln(\theta_i)} = \frac{\eta_i - 1}{\bar{\eta}} z_i \tag{21}$$

where $\bar{\eta} = \sum \eta_k z_k$ is the weighted average of η_k with the time shares as weights. Note that $\bar{\eta}$ is not a structural parameter and changes when the distribution of activities changes. Similarly, combining the $FOCs$ and differentiating with respect to T yields:

$$\frac{\partial ln(\mu)}{\partial ln(T)} = -\frac{1}{\bar{\eta}} \tag{22}$$

We can obtain the compensated demand elasticity of t_i by differentiating FOC_i and using 21:

$$\frac{\partial ln(t_i)}{\partial ln(1/\theta_i)}|_c \equiv \epsilon^c_{ii} = -(\eta_i - 1)(1 - \beta_i \cdot z_i) \tag{23}$$

where $\beta_i = \frac{\eta_i}{\bar{\eta}}$ is the Engel's curve parameter that captures how much demand for t_i changes when T increases. Activities with $\eta_i > \bar{\eta}$ expand and are thus "leisure luxuries," while demand for activities with $\eta_i < \bar{\eta}$ contract and are thus "leisure necessities."

The associated uncompensated demand elasticity is:

$$\frac{\partial ln(t_i)}{\partial ln(1/\theta_i)}|_u \equiv \epsilon^u_{ii} = -(\eta_i - 1)(1 - \beta_i \cdot z_i) - \beta_i \frac{ln(T)}{ln(1/\theta_i)} \tag{24}$$

We can use the envelope theorem to obtain two further useful results. Evaluating 19 at the optimal α_i^* and differentiating with respect to T and θ_i respectively yields the following results:

$$\begin{aligned} V_T &= \mu \\ V_T \cdot t_i &= V_{\theta_i} \cdot \theta_i \end{aligned} \tag{25}$$

where $V = v(\alpha_1^*, ..., \alpha_i^*)$ is the indirect utility function and $V_x = \frac{\partial V}{\partial x}$. The first result shows that the Lagrange multiplier is the shadow value of relaxing the constraint. While the latter result shows that increasing t_i and increasing θ_i have the same impact on utility.

Next we return to the upper nest optimisation. Evaluating eq 4 at $V(T, P)$ and taking first order conditions:

$$\begin{aligned} FOC_c: &\quad u_c = \tfrac{\Omega}{w} \\ FOC_T: &\quad u_v V_T = \Omega \end{aligned}$$

where Ω is the Lagrange multiplier on this upper nest problem. These FOCs pin down demand for T and c under standard assumptions on utility. To understand the beahviour of demand, log differentiate FOC_T with respect to T and to θ_i:

$$\begin{aligned} \frac{\partial ln(u_v)}{\partial ln(T)} + \frac{\partial ln(V_T)}{\partial ln(T)} &= \frac{\partial ln(\Omega)}{\partial ln(T)} \\ \frac{\partial ln(u_v)}{\partial ln(\theta_i)} + \frac{\partial ln(V_T)}{\partial ln(\theta_i)} &= \frac{\partial ln(\Omega)}{\partial ln(\theta_i)} \end{aligned} \tag{26}$$

The derivatives of $u_v(c^*, V(T,P))$ can be written as:

$$\begin{aligned} \frac{\partial ln(u_v)}{\partial ln(T)} &= \frac{T}{u_v}\left[u_{vv}V_T + u_{vc}\frac{\partial c}{\partial T}\right] \\ \frac{\partial ln(u_v)}{\partial ln(\theta_i)} &= \frac{\theta_i}{u_v}\left[u_{vv}V_{\theta_i} + u_{vc}\frac{\partial c}{\partial \theta_i}\right] \end{aligned} \tag{27}$$

To obtain expressions for $\frac{\partial c}{\partial T}$ and $\frac{\partial c}{\partial \theta_i}$, we need further assumptions about the shape of the utility function. Following Aguiar et al. (2021), we assume u_c is constant and can then differentiate FOC_c with respect to T and to θ_i and obtain:

$$\begin{aligned} 0 &= u_{cc}\frac{\partial c}{\partial T} + u_{vc}V_T \\ 0 &= u_{cc}\frac{\partial c}{\partial \theta_i} + u_{vc}V_{\theta_i} \end{aligned} \tag{28}$$

Combining 28, 27 and 26 yields:

$$\left[\frac{\partial ln(V_T)}{\partial ln(T)} - \frac{\partial ln(\Omega)}{\partial ln(T)}\right]\frac{V_{\theta_i}\theta_i}{V_T T} + \frac{\partial ln(V_T)}{\partial ln(\theta_i)} = \frac{\partial ln(\Omega)}{\partial ln(\theta_i)} \tag{29}$$

The term $\frac{\partial ln(\Omega)}{\partial ln(T)}$ is the change in leisure hours with the value of leisure and can be shown to be related to the inverse of the labor supply elasticity; one of the most studied parameters in labor economics. To see this note that the assumption that u_c is constant require $\frac{\partial ln(\Omega)}{\partial w} = 1$ and hence $\frac{\partial ln(\Omega)}{\partial ln(T)}|_{u_c} = \frac{\partial ln(w)}{\partial ln(T)}|_{u_c} = 1/e$, with e closely related to the "Frisch elasticity" of labor supply. Substituting 25 into equation 29 yields:

$$\left[\frac{\partial ln(\mu)}{\partial ln(T)} - \frac{1}{e}\right] z_i + \frac{\partial ln(\mu)}{\partial ln(\theta_i)} = \frac{\partial ln(\Omega)}{\partial ln(\theta_i)} \tag{30}$$

Combining this with 21 and 22 yields:

$$\frac{\partial ln(\Omega)}{\partial ln(\theta_i)} = \frac{\beta_i \cdot e - 1}{e} z_i \tag{31}$$

We have thus solved for the change in the value of leisure in terms of utility parameters z_i, β_i, e.

Using this fact, we can derive the change in leisure demand:

$$\frac{\partial ln(T)}{\partial ln(\theta_i)} = \frac{\partial ln(T)}{\partial ln(\Omega)}\frac{\partial ln(\Omega)}{\partial ln(\theta_i)} = e \cdot \frac{\beta_i \cdot e - 1}{e} z_i = (\beta_i \cdot e - 1) z_i \tag{32}$$

Leisure demand increases if $\beta_i \cdot e > 1$ and increases more for "leisure luxuries" with large values of β_i.

We can now return to the comparison of this non-homothetic case to the baseline case in 7 in the main text. The change in z_i is:

$$\frac{\partial ln(z_i)}{\partial ln(\theta_i)} = \frac{\partial ln(t_i)}{\partial ln(\theta_i)} - \frac{\partial ln(T)}{\partial ln(\theta_i)} \tag{33}$$

we can now derive the equivalent to the main result in 7 for the non-homothetic preference case. The relative change in T compared to z_i is:

$$\frac{\partial T/T}{\partial z_i} = \frac{\frac{\partial ln(T)}{\partial ln(\theta_i)}}{\frac{\partial ln(z_i)}{\partial ln(\theta_i)} \cdot z_i} = \frac{1}{(\frac{\partial ln(t_i)}{\partial ln(\theta_i)} / \frac{\partial ln(T)}{\partial ln(\theta_i)} - 1) z_i} \tag{34}$$

using 32 and 24, we can obtain the equivalent expression to 7, for the utility function used by Aguiar et al. (2021):

$$\frac{\partial T/T}{\partial z_i} = \frac{1}{(\frac{-\epsilon^c_{ii}}{(\beta_i \cdot e - 1) z_i} - (1 - \beta_i)) z_i} \tag{35}$$

note that when preferences are homothetic ($\beta_i = 1$) and $e - 1 = \eta_T$, this expression collapses to the case presented in the main text (7).[34] Otherwise, in the presence of non-homothetic preferences, there is an additional income effect that affects the elasticity of z_i. The magnitude of this effect depends on β_i; an entertainment innovation produces larger changes in z_i among leisure luxuries ($\beta_i > 1$) than among leisure necessities ($\beta_i < 1$).

We can now show how the presence of non-homothetic preferences affects the interpretation of empirical estimates on the LHS. Denote the empirical estimate of the LHS by ψ and re-arrange to obtain the solution for ϵ^c_{ii}:

$$\epsilon^c_{ii} = [1 + \psi \cdot (1 - \beta_i) z_i] \cdot \frac{\beta_i \cdot e - 1}{\psi} \tag{36}$$

For small values of ψ — as the case in our empirical application — the first term is ≈ 1 and

[34] Both e and η_T are closely linked to the "Frisch" elasticity of labor supply. However the link is derived under different assumptions. e assumes that u_c is unaffected by θ_i, e.g. with additive separable preferences, while η_T directly maps into the "Frisch elasticity" if there are no income effects. We therefore need the assumption about the values of e and η_T to make the two models coincide.

the main impact of β_i is to scale the denominator of the final expression up or down. For typical parameter estimates, this effect is modest relative to the gaps we identified in the text. Aguiar et al. (2021) estimate value for β_i ranging from 0.4 to 2.5 depending on the activity (Table 5). Allowing for $\beta_i \neq 1$ at most decreases the ϵ^c_{ii} estimate by factor 0.4, far less than the seven-fold gap between our estimate and previous calibrations. For the case of home entertainment, the impact of β even goes the opposite way and further widens the gap between our estimate and prior calibrations. Using the parameter value of Aguiar et al. (2021) ($\beta_i = 2.4$) increases the estimate of ϵ^c_{ii} by close to 2.4 times.

8.3 Individual Level Labor Supply Model

We now present a simple framework to model the labor supply decision of an agent that experiences an improvement in an entertainment technology. This approach provides predictions about extensive and intensive margin responses. We will focus on the retirement implications of the launch of television and show that it provides a set of predictions which help distinguish the impact of a technical change that enhances the value of leisure from other potential changes that do not predominantly affect retirement age cohorts.

Consider an individual with preferences over leisure $(1 - l)$ and consumption (c) and utility function $U(c, \xi(a)l)$, with $\xi(a)$ capturing heterogeneity in the value of leisure in the population. We follow the retirement literature and assume that work becoming more taxing as people age (e.g., French (2005); French and Jones (2011)) and assume that $\xi(a)$ is an increasing function of age, denoted by $\beta(a)$ with $\beta'(a) > 0$, and a shock ν that is independent of age: $\xi(a) = \beta(a) + \nu$. Following Lazear (1986), working is also associated with a fixed cost of work denoted by x .

A leisure-enhancing technology that increases ν affects work hours at the intensive margin (l^*) and the retirement age at the extensive margin $(\tilde{a})$.

Assume the utility function is quasi-linear with $U(c, \xi(a)l) = c - \frac{\xi(a)}{1+1/\epsilon}(\frac{l}{\xi(a)})^{1+1/\epsilon}$, with ϵ representing the labor supply elasticity. And the budget constraint when working is $c = w \cdot l - x + b_0$ and $c = b_0$ when not working, with non-wage income b_0. Following Lazear (1986), a worker pays a fixed cost of working x, which has the effect that working a small number of hours is undesirable and workers will either work substantial hours or not at all:

$$max\ U(c, \xi(a)l) \tag{37}$$

$$\text{s.t. } c = \begin{cases} w \cdot l - x + b_0 & l \leq 1 \\ b_0 & l = 1 \end{cases}$$

Denote the value of leisure of a person is just indifferent between working an no by $\xi(\tilde{a})$. Figure 6 illustrates this case. All people with $\xi(a) > \xi(\tilde{a})$ will not work and people with $\xi(a) < \xi(\tilde{a})$ will work, implying that people with age $a > \tilde{a}$ are retired.

Using this definition, we can derive the retirement age in this economy. The marginal retiree is indifferent between working and not working. The utility when not working is $U_0 = b_0$ and equals the utility at the interior point $U_0 = U^*$. Utility at the interior solution (U^*) follows from utility maximization. At an interior solution the first order conditions imply that $l^* = \xi(\tilde{a}) \cdot w^{\epsilon}$ and hence $U^* = b_0 - x + \frac{w^{1+\epsilon}}{1+\epsilon}\xi(\tilde{a})$. Combining this result with $U_0 = U^*$, we get an implicit expression for $\tilde{a}$:

$$\xi(\tilde{a}) = \frac{x(1+\epsilon)}{w^{1+\epsilon}} \tag{38}$$

We can use this expression to derive comparative statics and analyze the impact of leisure-enhancing technologies. Such technologies increase ν and have two effects on labor supply. First, they affect the optimal labor supply:

$$\frac{\partial l^*}{\partial \nu} = w^{\epsilon} > 0$$

For all workers at an interior solution, leisure consumption increases by w^{ϵ}. The greater utility of leisure leads to a marginal reduction in work hours.

Moreover, such technological changes have extensive margin effects and push a greater share of people to shift from l^* to $l = 0$. The effect operates through a falling retirement age. Using the implicit function theorem on equation 38 yields:

$$\frac{\partial \tilde{a}}{\partial \nu} = -\frac{1}{\beta'(\tilde{a})} < 0$$

A rising value of leisure thus leads to earlier retirements and increased exit from the labor force. Figure 6 shows the intuition behind this result. The rising value of β_0 pivots the indifference curve upward and makes it steeper. This implies that the new marginal retiree has $\xi(\tilde{a}^{'}) < \xi(\tilde{a})$, and hence $\tilde{a}' < \tilde{a}$. The new marginal retiree is thus younger, and individuals with age between $\tilde{a}'$ and $\tilde{a}$ will have exited the labor force.

The model offers three simple insights. First, leisure-enhancing technologies reduce labor supply both at the extensive and intensive margins. Second, the group that responds most are older workers whose relative value of leisure is highest. This group is at the margin of labor force participation to begin with and therefore most likely to respond to leisure-enhancing technologies by exiting the labor force. Third, while the value of leisure changes only marginally, the labor supply responses are still substantial among some groups. A fixed cost of work implies that some people jump from near full-time participation to not working at all.

While the intensive effects and the simplicity of the results hinge on the functional form assumption, the extensive margin effects hold for a broad set of utility functions. These participation predictions are one of the few general predictions of the labor supply framework that hold independently of the parametric assumptions about the utility function.

9 Appendix C: Measuring TV Access

9.1 ITM Model Details

We use an ITM model to compute the TV station variable. We first digitize data about stations addresses, technical features (height, power, frequency), and launch dates from Television Factbooks. We then feed these data, together with terrain data, into the ITM model and calculate the signal quality of each station at the centroids of all US mainland counties. We define catchment areas using a 90-90 reliability metric (90 percent reliability with 90 percent confidence) and set the cutoff signal quality at -50. This yields information about TV stations at the county level. Finally, we aggregated the data to the level of our outcome data by taking population-weighted averages over counties.

The catchment areas of the stations are updated annually, as changes in broadcast technology alter the catchment areas over time. In each year, we base our calculations on the station details in September of that year, when most of our raw data sources are published.

9.2 Measurement Error in the DMA Data

Gentzkow (2006) approximates 1950's broadcast ranges with Nielsen media markets, or Designated Market Areas (DMAs), that are based on 2003 viewership. A DMA is a group of counties around a metropolitan area. The approximation takes the year in which the first station in a DMA began operation and assumes that each county in that DMA received a signal in that year. We found that 1960's coverage maps show differences between historical broadcast ranges and the 2003 DMAs. The DMA approximation sometimes underestimates and sometimes overestimates how far signals reached. The next two subsections give examples of each case. These are not representative, as we chose them specifically for exposition of the two types of problems with the DMA approximation.

9.2.1 An Example of DMA Underestimation (A type II error)

Proximal cities confound the DMA approximation of TV access. For example, panel (A) of figure 10 shows a coverage map of Kansas City from the 1967 *TV Factbook*. The blue line is the broadcast ring as defined by those counties that have over 50 percent coverage according to the map. Panel (B) overlays in red the Kansas City DMA. The DMA is too small–it excludes counties to the northwest that were likely covered. Moreover, for a region with little variation in terrain, the irregular shape of the DMA suggests that it cannot reflect the roughly circular true broadcast range.[35]

Let TVYEAR_i denote the year in which county i first had TV access. In panel (B), the DMA approximation assigns the highlighted counties between the two rings a TVYEAR of 1954. However, those counties fall well within the range of the Kansas City tower, and that tower started broadcasting in 1950. Therefore the true TVYEAR of the highlighted counties is likely 1950, not 1954. This misclassification owes to the nearby DMAs, Topeka and St. Joseph, whose broadcasts began in 1954. While it is true *today* that the highlighted counties are closest to the Topeka and St. Joseph signals, and are therefore not in the 2003 Kansas City DMA, those counties are close enough to Kansas City to have viewed Kansas City broadcasts in 1950.

The TV ownership data from Gentzkow and Shapiro (2008) confirm that this is a case in which today's DMAs do not align with 1950's signals. The DMA data assign the highlighted counties in panel (B) as not receiving a TV signal until 1954, four years after the counties in the red Kansas City ring. If that were true, we ought to observe the highlighted counties buying TVs well after the Kansas City counties. Panel (A) of figure shows that in fact the timing of TV purchases is almost identical across the two groups, consistent with the hypothesis that Topeka and St. Joseph viewers received a 1950 signal from Kansas City. Substantial TV ownership in a county before that county's DMA-approximated TVYEAR is evidence of measurement error arising from signal overlap.

When signals overlap like this, DMAs underestimate coverage. The overlap between Kansas City and Topeka, for example, leads the DMA data to underestimate how many counties the Kansas City broadcast reached in the 1950's. Spot-checking coverage maps suggests that DMAs can also overestimate coverage.

9.2.2 An Example of DMA Overestimation (A type I error)

Today's DMAs sometimes extend further from city centers than historical signals did. Panel (C) of figure 10 shows a *Factbook* coverage map of Minneapolis-St. Paul. The blue line rings counties whose coverage exceeded 50 percent. Panel (D) adds the Minneapolis-St. Paul DMA in red. That DMA is too large, in that it includes the highlighted counties that were likely out of reach of the broadcast, which leads to overestimation of coverage. The highlighted counties have a

[35] For two reasons, the *Factbook* maps ought to be taken only as suggestive regarding true 1950's signal reach. The first is that these maps were not published until the 1960's, and tower technology–power, height, etc.–improved substantially over time. The second is that the shading in the maps reflects surveys of viewership, not measures of signal strength. County coverage exceeding 50 percent for a station means that over 50 percent of households in the county watched that station. Our measurement of signal reach will not rely on these maps.

DMA TVYEAR of 1948, since that is when the first Minneapolis station began operation. But many of those counties appear to be too far away from the tower to receive the early Minneapolis signals. Panel (B) of figure shows that TV purchases in the highlighted counties–the group inside the DMA but outside the mapped broadcast range–lagged purchases in the counties inside the *Factbook* coverage area, consistent with the hypothesis that the DMA overestimates 1950's signal reach. That pattern remains after controlling for county characteristics like income and population that are associated with TV ownership.

9.2.3 Causes and Prevalence of Measurement Error

This section moves beyond examples to the causes of measurement error and evidence on the prevalence of those causes. To start with underestimation, the two conditions under which the signal overlap problem arises are: Neighboring DMA towers (1) are close enough for signals to overlap and (2) started broadcasts in different years[36]. The closer the towers and the further apart the initial broadcast years, the larger the potential measurement error. To find possible areas of overlap, we ranked pairs of DMAs by their distance apart. There are 166 unique pairs of DMAs whose towers are less than 100 miles apart (a typical broadcast radius) with broadcasts beginning in different years. Among them are the Kansas City, Topeka, and St. Joseph pairs. Other metropolitan areas such as Pittsburgh and Cleveland are close enough to smaller neighboring stations like Youngstown to create the same overlap issue.

Overestimation, by contrast, can arise because of improvements in TV towers over time. In most cities, the 1950's saw expanded broadcast ranges through both upgrades to existing stations and also construction of new towers. The 2003 DMAs are therefore prone to overstate early 1950's signal reach, when towers were weaker. As shown in figure 12, the average height above ground of a commercial tower in 1948 was 483 feet, and already by 1960 that had increased to 629 feet. Some stations moved to higher ground, and tower height above average surrounding terrain rose from 721 to 992 feet. Average visual power jumped from 19 to 170 kilowatts over that period, and average aurul power increased from 11 to 87 kilowatts.[37] The fixed DMAs do not capture shifts in broadcast areas that followed changes in tower technology.

These measurement issues tend to affect particular types of counties. The DMA approximation always gets major cities right. Underestimation and overestimation occur at the fringe of the broadcast areas of those cities, as the figure 10 examples show with Kansas City and Minneapolis-St. Paul, and the fringe plays a key role in estimating TV's effects. Gentzkow (2006) exploits

[36]Condition (2) is necessary because if two towers were close but started broadcasts in the same year, then all surrounding counties would get a signal in the same year, so proximity alone would not lead to misclassification. Terrain also matters–mountains could prevent overlap–and our measurement of TV access will account for variation in elevation.

[37]Power does not map directly to broadcast reach, as higher frequency channels require more power to operate.

broadcast rings to identify the causal effects of TV on voter turnout. The idea is that since TV reception reached about 100 miles from a broadcast tower, counties just inside and outside of that radius comprise treatment and control groups. Using this method, variation in access to TV is "driven by whether a county happened to fall within the roughly 100-mile radius of television broadcasts" (p. 945), so measuring that radius accurately is especially important for inference.

We took the evidence presented thus far as reason to pursue a more precise measure of TV access. Those measurements, constructed using digitized *TV Factbook* data and the Irregular Terrain Model (ITM) of signal propogation are discussed in section 2.1 of the main text. To validate the ITM measurements, we turn next to comparisons of key findings in the literature using the DMA approximation and ITM data.

10 Appendix B: Empirical Appendix

10.1 Social Security Sample

The Social Security Act of 1935 introduced Federal Old Age Insurance in the United States. Individuals over the age of 65 received benefits, and payments were based on contributions people made across their work histories. To keep track of individual contributions, the Social Security Administration (SSA) started recording individual earnings data in 1937. Initially this covered all wage and salary workers (excluding railroad workers) under age 65 who were employed in the private sector in the U.S. and Alaska and Hawaii, which were then territories (Long, 1988). From the outset, the system thus covered a substantial share of the U.S. workforce; in 1937 it was estimated that around 32 million workers, or roughly 60% of the labor force, were covered (Wasserman and Arnold, 1939). Workers not excluded from the system included certain non-covered occupations (e.g. the self-employed), workers aged 65-74, and the unemployed or workers in unemployment relief programs. Coverage was expanded over the following decades, with major expansions in 1951, 1954 and 1956. The expansions broadly affected workers in four categories: government employees, the self-employed, military personal, and agricultural workers. To work with a consistent sample, we drop occupations that first receive coverage during this period.[38] Since the data only report occupation and industry in 1977, we also exclude individuals that first appear in the earnings records in one of the three extension years in the 1950's and are older than 30.[39]

At the beginning of the sample, the Social Security system excluded the following groups:

[38]This excludes 3,714 individuals. We exclude workers in occupation groups: 42, 43, 44, 36, 10, 11, 7; in occupations: 821, 822, 980, 981, 982, 983, 984, 824; in major industry group: 11; industry group: 48, 49, 50, 51; industries: 927, 937, 769; and workers in areas with a farming to population ratio over 10%. Additionally, we exclude veterans who appear in the data in 1957.

[39]This drops an additional 1,996 individuals.

"agricultural employment, work for Federal, State and local governments, employment by certain non profit organizations or institutions, railroad employment, domestic service in private homes, and all types of self employment." Moreover, workers over the age of 65 did not contribute to Social Security in 1937 and 1938 and their employment was not recorded (Social Security Bulletin, Vol. 70, No. 3, 2010), so we set employment to missing for these cases. In 1951 the self-employed (except members of professional groups), farm laborers and domestic workers were included in the system. Additionally, worker in nonprofit organization could join the system if they received at least \$100 in pay during the calendar year. Reforms broadened coverage further in 1955. These reforms relaxed restrictions on farm workers, the self-employed and expanded the scope for voluntary participation of state and local government employees. Farm laborers were included if they passed a "cash-pay" or "regularity-of-employment" test. This required a cash income over \$150 from a single employer, or employment on a time basis of at least 20 days with a single employer. Finally, in 1956 soldiers on active duty, previously excluded self-employed professions and optionally police and firefighters in state and local retirement systems became covered. To avoid individuals dropping in and out of employment due to changes in the earning threshold, we code all workers as employed if they earn over \$50 and non-employed if earnings are below \$50.

10.2 Summary Statistics

Our baseline sample comprises of 325,130 person-year observation, 31,653 individuals and spans 134 local areas. As described above, these areas split the mainland U.S. into MSAs and rest of state areas. We present summary statistics of our sample in Table 5. A few observations are worth highlighting. First, the SSA employment measures are not directly comparable with variables from the Census. The previous section describes how the SSA defined employment and we use this definition. Also note that using SSA employment definitions has become a common practice in a sizable literature that analysis the U.S. labor market with administrative records. The picture is broadly consistent with Census data and we discuss employment trends more below. Second, it is worth exploring the representativnes of the sample. While a representative sample is not necessary for the validity of the analysis, understanding the sample helps understand the summary statistics. Our sample is based on the 1978 CPS and thus becomes less representative of the U.S. population as we go further back in time. In particular, groups with higher mortality or migration rates are underrepresented. As a result, the sample includes somewhat fewer men (41% instead of 49%) and minority workers (9% instead of 10%) and is younger (38 instead of 44) than the U.S. population of the time. All in all, the sample is reasonably close to the aggregate U.S. population. A major strength of the experiment is that it touches broad range of society and we can measure heterogeneous effects by sub-groups and strengthen the external validity of our results. For instance, the effect of television may look differently in a population with a different demographic make-up. Below we explore this formally and re-weight our sample to obtain the average treatment effect for the U.S. society.

Reassuringly, the CPS-SSA show the familiar life-cycle labor supply pattern and closely align with the familiar 1950 US Census patterns. The employment to population rate for men follows a U-shaped patterns: it rises until age 30, then plateaus and starts declining from age 50. For women, employment rates start at a lower level and decline during the child bearing years, then recover somewhat in the late 30s until they start declining later in life. These patterns are well known and are broadly consistent with the analysis in Mcgrattan and Rogerson (2004). The SSA data thus appears to capture the core feature of the life-cycle labor market accurately.

Our data show a rapid increase in retirement rates in the 1950's. Figure 7 shows that the retirement rates for over 65 year olds almost doubled from around 30% to nearly 60%. Our measure of retirement differs somewhat from Census definitions of labor market activity. We define retirement as a permanent with-drawl from the labor force, as measured by Social Security contributions. Census measures typically focus on employment in one specific reference week. These definitions make a difference to the level but not the trend in inactivity, both series show a sharp decline in labor market activity among the over 65 year olds during the 1950's. A second striking feature of Figure 7 is the rise in retirement among "younger" cohorts. Retirement is less common among people aged between 50 and 65 but the trend in the 1950's clearly points upwards too. Retirement rates among these "younger" workers almost doubled in the 1950's. This trend is particularly remarkable because these age groups are typically not eligible for Social Security, which suggests that other factors beyond social insurance played a role in growing retirement trends.

Finally, we provide additional detail on the variation from the television rollout. Figure 9 shows the time series aspect of the rollout. At the start of the license freeze in 1950 substantial differences existed across the U.S.. Multiple stations were already available in a few early adopting locations but most Americans had only limited exposure to television. This changes with the lift of the license freeze in 1952. In the following two years television spread throughout the country. The figure illustrates that much of the variation in the television rollout over time is down to the license freeze "accident," which helps our identification strategy.

10.3 Robustness Checks

10.3.1 Stocks vs Flows

The baseline regressions use employment flows as outcome ($\Delta Employment$). This section shows alternative results with stocks as outcomes and discusses how the two relate to each other.

The stock regression assumes that the time invariant individual effect absorbs potential selection problems and, conditional on these FE, the treatment assignment is random. We show results using this specification in Column 1 of table 6. The pattern is similar to the main results of the paper, we find significant negative effect, particular on older workers. The magnitude of these estimates is slightly larger, the effect on younger workers is now marginally significant and we find significantly smaller effects on women. The identifying assumption of this basic stock regression is violated when the individual effect vary over time, for example because a persons circumstances or norms change. A simple way to account for such changes is to include a lag of the dependent variable (LDV). Column 2 does so and finds largely similar results. The only difference is that the coefficient on women is smaller (potentially becuase shifting norms around female labor force participation weren't fully cpatured by the time invariant fixed effect). Adding a LDV, however, creates it's own econometric problems. Specifically, the coefficient on the LDV is biased downward, an issue known as "Nickell bias." Arellano-Bond developed an estimator to fix this bias and suggest instrumenting the LDV with further lags of the same variable. Column 3 implements the Arellano-Bond estimator (AB), using an instrument that lags the LDV by one additional period, and finds similar results. The credibility of these estimates, however, hinges on strong timing assumptions and it is unclear whether the AB estimator yields an unbiased estimate.

The final specification is the baseline specification used in the paper and focuses on flows directly. One way of thinking of this regression is as an extension of the LDV models that imposes a coefficient of 1 on the LDV.

10.3.2 LATE vs ATT: sample weights

Our SSA-CPS data follows the 1978 CPS cohort throughout their life. The sample is representative of the 1978 population but becomes less representative as we go back in time. The lack of representativeness does not cause problems for the internal validity of the results, but it does limit the external validity. Specifically, our estimate will suffer from "survivor bias" if people with the biggest response are more likely to die and thus less likely to appear in our data. We don't think this is particularly likely, but if it is the case our LATE estimate is a lower bound for the Average Treatment Effect on the Treated (ATT). To get a better sense of the importance of such survivor bias, we re-weighting our sample putting more weight on groups with small survival rates. Our target population is the population of the 1950 and 1960 U.S. population Census and we weight our observations to match those population totals.[40] Specifically, we target population aggregates in an MSA, as well as their education and age demographics.

Table 7 shows the baseline results with the weighted sample. The main takeway is that the results are broadly similar to those reported in our baseline results (see Table 1). The point estimates

[40] We linearly interpolate values in between the 1950 and 1960 Census.

are slightly larger for both groups and now also boarder line significant for the younger age group.

10.3.3 Effect Heterogeneity

We here analyze heterogeneity in the response across demographic groups. The first column allows for different effects among more mobile individuals. Mobile individuals are more likely to leave the fixed MSA that we assign them to during the analysis and by testing treatment effects on this sub-group, we can assess how much such moves may attenuate the results. We define a dummy for high vs low mobility individuals and look at differences in the effects. We do not have data on moves in the 1950's and instead use the CPS migration supplement to infer moving propensity. We classify people as mobile if the moved out of MSA between 1975 and 1976 and test how much mobility attenuates results. The difference in effects is insignificant and quantitatively small (Table 8, column 1). This suggests that the attenuation bias from mobility is relatively minor.

The next columns show heterogeneity cuts for other demographic groups. Column 2 looks at differences by marital status and column 3 by education level, finally column 4 looks at the effect on minority workers. None of these groups show significantly different responses.

10.3.4 Migration and Intention to Treat

The baseline estimates treat place of residence as fixed and estimate intent-to-treat (ITT) effects. This appendix explores how these ITT effect relate to the local average treatment effect. Generally, migration could have two potential effects on the results. First, endogenous moves towards television could lead to selection effects, second random moves will lead to mis-measurement of television exposure. The first issue, selection effects, are taken care of by the individual fixed effects in our analysis. The focus of this section is instead on the second problem, which we call the imperfect compliance challenge, in the spirit of ITT effects. The standard approach in the literature is to divide the ITT estimates by the rate of compliance. In our setting, the denominator would be the fraction of people who migrate outside the treatment area. We additionally require information on the treatment effect in the non-complier population. In a set-up with a binary treatment non-compliers don't access the treatment and have a zero treatment effect. However, with multiple treatment dosages, non-compliers may still experience some treatment effects. The relation of the ITT and ATT can be expressed as: $ITT = ATT \times \sigma + NCTE \times (1 - \sigma)$. Where σ is the compliance rate, or the share of people who lived in a different MSA than we observe, and $NCTE$ is the treatment effect experienced by these non-compliers. Note that with a binary treatment $NCTE = 0$ and the ATT becomes the familiar IV estimate that scales the ITT up by the compliance rate: $ATT = ITT/\sigma$.

We first calculate the approximate level of non-compliance in our sample (σ). This requires data on migration patterns. The CPS-SSA linked data only includes imperfect information on these rates and we use the matched 1978 CPS migration supplement to estimate migration rates. Many people move every year, but only a small fraction of these moves affects our results. In particular, only moves that cross MSA boundaries are relevant. According to the 1978 CPS migration supplement, 5% of our sample left an MSA during the three year window 1975-1978. This group are clearly non-compliers and we can use this group for a benchmark exercise with $\sigma = 0.95$. To calculate the ATT we also need an estimate of the $NCTE$ and Table 8 reports treatment effects for this non-complier group in column 1. Using $\sigma = 0.95$ and $NCTE = -0.224$ in the ATT formula yields an ATT of -0.273, very close to the ITT estimate of -0.270.

The previous estimate is likely a lower bound for the true ATT as it only takes migration between 1975 and 1978 into account. The share of people who left the MSA in the 18 year window from our sample period to the 1978 CPS is larger. If we assume stationary migration rates, we can extrapolate the 18 year rate as: $\sigma = 0.05 + \sum_{t=1}^{5} 0.05(1-p)^t$, where p is the rate of repeat migration. A high value of p implies that some people are intrinsically more mobile and move frequently. We use panel data from the NLSY79 to get a sense of these repeat migration rates and find rates around $p = 0.3$. This implies $\sigma = 0.15$ and together with our previous $NCTE$ estimate yields an ATT of -0.278, again similar to the baseline estimates. To push this to an extreme, assume next that people only move ones ($p = 0$). In this scenario the ATT=-0.291, and therefore still in the same ballpark as our baseline estimates. This is of course an unrealistic assumption but illustrates that the results are reasonably robust to alternative assumptions about migration patterns.

11 Appendix Figures and Tables

Figure 6: Marginal Retiree

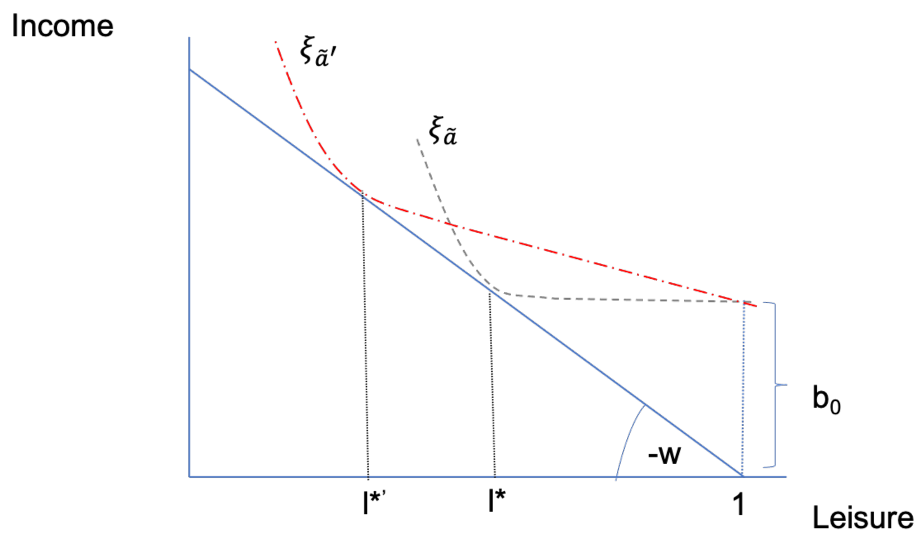

Notes: The figure shows the indifference curve of the marginal retiree, a person who is just indifferent between working and not. The age of the marginal retiree is indicated by $\tilde{a}$. The dashed line is a case with low β_0 and the dash-dot line is a case with higher β_0.

Figure 7: Retirement Rates

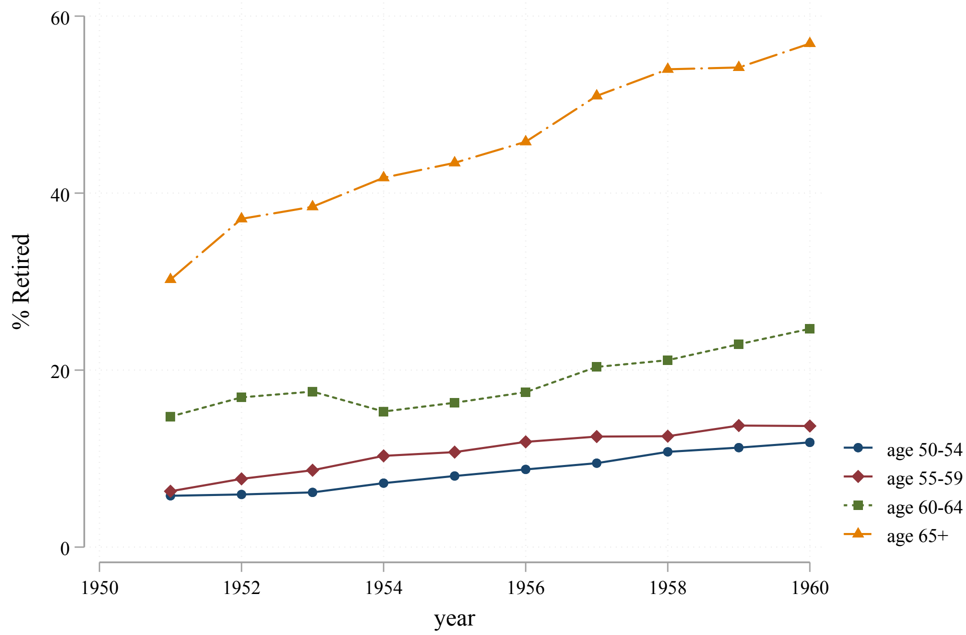

Notes: The figure shows retirement rates among older workers during the 1950's. Retirement is defined as no observed employment in the Social Security records until the end of our sample (1978). Source: linked SSA-CPS data.

Figure 8: Employment-to-Population Rates over the Life-Cycle

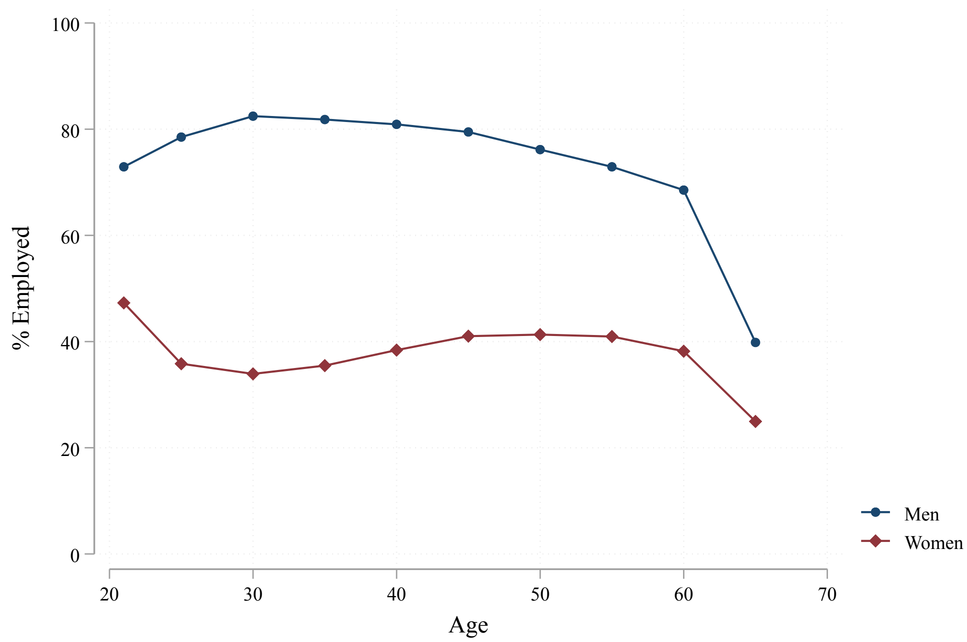

Notes: The figure shows employment rates by age and gender. Each dot shows the average for a five year age window, averaging employment rates over the full sample period. The first and last bins respectively show averages for the age groups 21-24 years and 65+. Source: linked SSA-CPS data.

Figure 9: Number of Stations Available Over Time

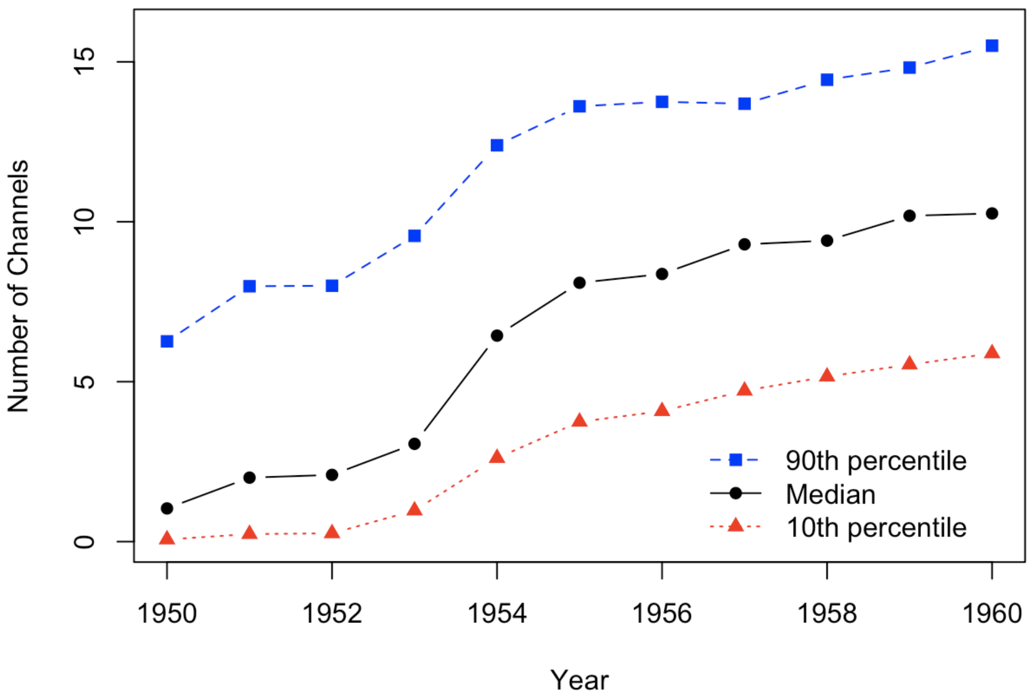

Notes: The figure shows the number of television stations in the U.S. between 1950 and 1960. It shows this for a median person, as well as at the 90th and 10th percentile of the distribution.

Figure 10: Coverage Maps and Designated Market Areas

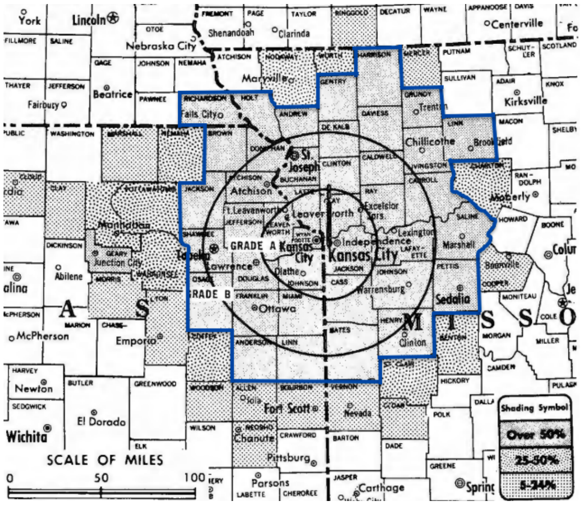

(A) Kansas City coverage map ring (in blue)

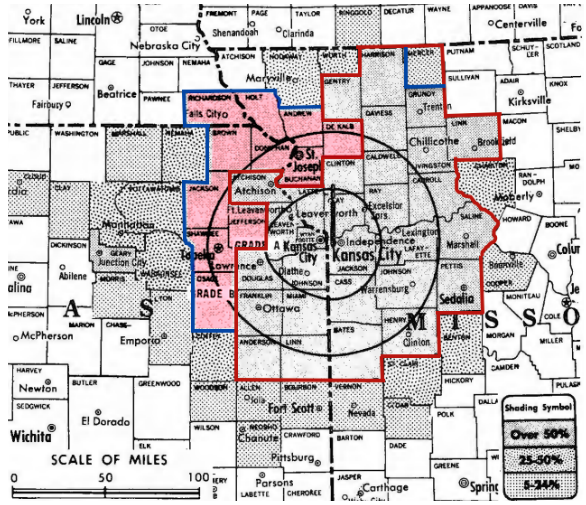

(B) Kansas City DMA ring (in red)

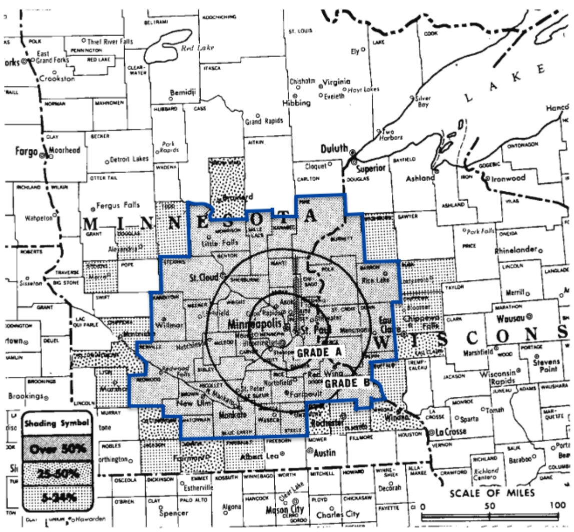

(C) Minneapolis coverage map ring (in blue)

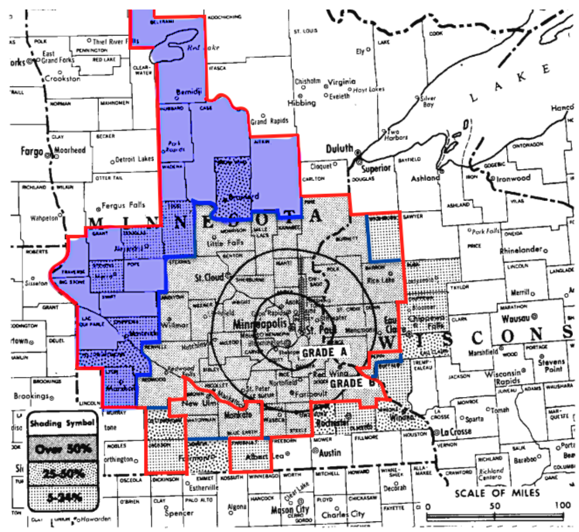

(D) Minneapolis DMA ring (in red)

Figure 11: TV Purchases Patterns

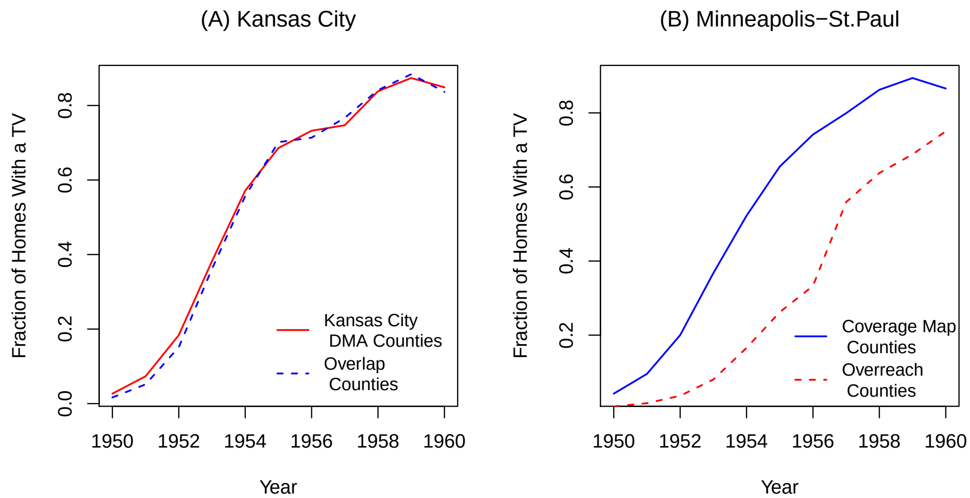

Notes: Panel (A) shows average TV ownership around Kansas City for counties in the groups indicated in the legend. "Overlap Counties" refers to those highlighted in Figure 10. In Panel (B), for Minneapolis-St. Paul, "Coverage Map Counties" refers to those ringed in Figure 10 (C), whose coverage exceeds 50 percent according to *TV Factbook* coverage maps. "Overreach Counties" refers to those highlighted in Figure 10, which fall inside the Minneapolis-St. Paul DMA but outside the *TV Factbook* broadcast range.

Figure 12: Broadcast Technology Improvements

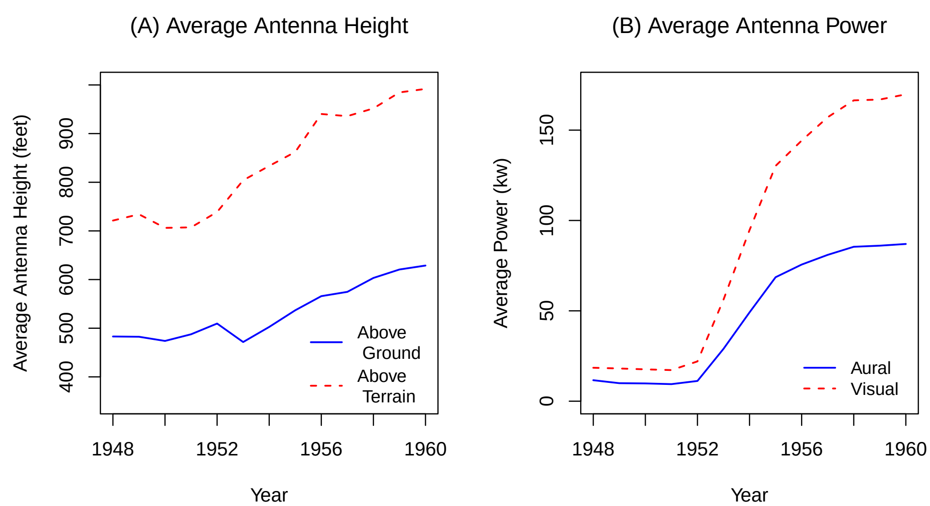

Notes: The figure shows the increases in broadcast tower height and power over time. Data are digitized from the *TV Factbook*, as discussed in the main text.

Table 5: Summary Statistics

	Observation	Average	s.d.	Min	Max
Employed	292,448	56.32	49.6	0	100
Quarters worked	292,376	2.02	1.9	0	4
TV Stations	292,448	7.98	4.3	0	16.65
Years of schooling	292,448	11.8	3.37	1	19
High school graduate	292,448	0.54	0.498	0	1
Year of birth	292,448	1917	10.91	1881	1937
Female	292,448	0.59	0.49	0	1
Non-white	292,448	0.08	0.27	0	1
Recent move	288,942	0.05	0.22	0	1
Ever married	292,448	0.95	0.21	0	1

Notes: The table reports summary statistics for the SSA-CPS sample. Employment and age information is based on SSA records and spans the years 1946-1960. The data is annual from 1951 to 1960 and uses multi-year averages for the periods 1937-1946 and 1947-1950. We restrict the sample to adults (over age 21 at the time). Data on gender, marriage, mobility, race and schooling is based on linked 1978 CPS records. Data on TV stations is computed using records from digitized Television Factbooks in an ITM signal propagation model.

Table 6: Effect of Television on Employment Stocks and Flows

	(1)	(2)	(3)	(4)
		Employment		ΔEmp
Effect of TV stations ×				
Age over 50	-0.572***	-0.462***	-0.500***	-0.276***
	(0.155)	(0.102)	(0.109)	(0.0683)
Age under 50	-0.270**	-0.203**	-0.199**	-0.0911
	(0.127)	(0.0856)	(0.0933)	(0.0590)
Female	0.377**	0.127	0.157	0.0250
	(0.163)	(0.109)	(0.118)	(0.0831)
Observations	292,448	292,448	260,895	292,448
R-squared	0.736	0.774	0.146	0.048
Year FE	Yes	Yes	Yes	Yes
Person FE	Yes	Yes	Yes	Yes
Specification	Stock	LDV	AB	Flow
Mean DV age over 50	51.89	52.01	52.24	0.337
Mean DV age under 50	55	56.11	56.37	1.421

Notes: The table reports effects of television stations with alternative specifications. All regressions control for person and sex by year FE. Specification ”Stock” uses an employment indicator as outcome, ”LDV” controls for the laged dependent variable, ”AB” implements the Arellano-Bond estimator, instrumenting the lagged dependent variable with one further lag, ”Flow” uses changes in employment as outcome. Other specification details are as in 1, column 1.

Table 7: Individual-level Effects of TV on Employment - Weighted Sample

	(1)	(2)	(3)	(4)	(5)
			$\Delta Employment$		
Effect of TV stations $\times$					
Age over 50	-0.316***	-0.328***	-0.327***	-0.342***	-0.324***
	(0.111)	(0.111)	(0.118)	(0.124)	(0.118)
Age under 50	-0.284***	-0.282***	-0.248**	-0.258**	-0.247**
	(0.101)	(0.101)	(0.109)	(0.115)	(0.110)
Female	0.0575	0.0574	0.0512	0.0629	0.0515
	(0.159)	(0.159)	(0.158)	(0.186)	(0.158)
Observations	288,266	288,266	288,266	288,266	288,266
R-squared	0.073	0.073	0.048	0.073	0.073
Year FE	Yes	Yes	Yes	Region X Year	Yes
Person FE	Yes	Yes	Yes	Yes	Yes
Age FE	No	Yes	Yes	Yes	Yes
Quartic Trends	No	No	Region	No	Demographics
Mean DV age over 50	0.350	0.350	0.350	0.350	0.350
Mean DV age under 50	1.421	1.421	1.421	1.421	1.421

Notes: The table replicates Table 1 and additionally uses sample weights. $***p < 0.01, **p < 0.05, *p < 0.1$

Table 8: Heterogeneous Effects of TV on Individuals Aged 50+ by Demographic Groups

	(1)	(2)	(3)	(4)
Stations	-0.270***	-0.229***	-0.258***	-0.277***
	(0.0674)	(0.0778)	(0.0702)	(0.0684)
Stations $\times$ $\mathbb{1}$(Mobile person)	-0.0460			
	(0.144)			
Stations $\times$ $\mathbb{1}$(Ever married)		-0.0503		
		(0.0535)		
Stations $\times$ $\mathbb{1}$(High school graduate)			-0.0732	
			(0.0452)	
Stations $\times$ $\mathbb{1}$(Some college or more)			-0.0372	
			(0.0623)	
Stations $\times$ $\mathbb{1}$(Minority)				0.0204
				(0.0758)
Observations	288,942	292,448	292,448	292,448
R-squared	0.047	0.048	0.048	0.048

Notes: The table shows regressions of employment on available TV stations with interactions for the listed demographic groups. The specification is the baseline specification in column 3 of Table 1. Mobile person: person moved MSA between 1975 and 1976. The omitted category in column 3 is "High school dropout." Robust standard errors in parentheses. $***p < 0.01, **p < 0.05, *p < 0.1$