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

George Fenton and Felix Koenig†

June 17, 2023

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. A TV station launch reduced the probability of working by around 0.3 percentage points driven mainly by an increase in retirement rates among older age groups. We use a representative agent model to assess the impact of entertainment innovations on aggregate labor supply trends more broadly and find that TV had a substantial impact, while subsequent technologies had small effects.

*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: Center on Budget and Policy Priorities (gmfenton@umich.edu); 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 0.35

[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).

percentage points (0.5 for men, 0.2 for women). 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 markedly. These exits are predominantly into retirement. Retirement rates among older workers increase by 0.3 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 education, gender, racial, and marital 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 socks 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

[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.

implies a large improvement in the value of this activity. Since time shares are observable, we 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. We first show that an hour of TV crowds out about three minutes of work time and fifty-seven minutes of other leisure activities, meaning people mainly substitute TV for other leisure activities. In the model, the fact that people mostly give up other leisure activities to watch television implies 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. Other entertainment innovations with smaller impacts on American time use thus had only small effects on aggregate labor supply.

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); Baker and George (2010); Campante and Hojman (2013); Thomas (2019); Kearney and Levine (2019); Kim (2020); La Ferrara et al. (2012); Angelucci et al. (2021) and Chadi and Hoffmann (2021)).

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 Entertainment Technology and Labor Supply

Two types of frameworks are widely used to model labor supply behavior: a micro and a macro approach. Aguiar et al. (2021) use the macro approach and model aggregate labor supply trends (i.e. total work hours) with a representative agent framework. The micro approach, instead, models individual decisions and provides more flexibility to study heterogeneity in behaviour (and naturally distinguishes extensive and intensive margin effects). We present both types of frameworks and derive the implications of entertainment technologies. The macro approach most closely mirrors prior work in the area, while the micro approach provides more nuanced predictions that can help distinguish the value of leisure mechanism from other potential mechanisms.

2.1 The Representative Agent Model

First, consider a representative agent framework, similar to the one used in Aguiar et al. (2021) to study the impact of video games.[7] This model studies changes in aggregate labor supply and thus combines the intensive and extensive responses. A clear advantage of such a representative agent framework is that it offers a parsimonious way of thinking about aggregate employment trends, without having to track corner solutions.

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

$$u(c, v(\alpha_1, ..., \alpha_n)) \tag{1}$$

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 9.2, finding small effects on the predictions we study. We therefore focus on the simpler homogenous utility case in the main text.

The utility maximisation problem of this agent is:

$$\max_{a_i, c} \quad u(c, v(\alpha_1, ..., \alpha_n)) \quad \text{s.t. } L + T \leq 1, \sum_{i=1}^{n} t_i \leq T, c \leq w \cdot L, \alpha_i = \theta_i \cdot t_i \tag{2}$$

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.[8] 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.[9] In addition, it will be useful to define the compensated demand elasticity for activity i: $\frac{\partial ln(\alpha_i)}{\partial ln(\theta_i)} \equiv \epsilon_{ii}^c - 1$, 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.

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

[8] 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.

[9] 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 $\frac{-T}{L}$: $\eta_L = \eta_T \frac{-T}{L}$.

Solving the utility optimisation, we can derive the following two results for the impact of 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 8 for derivations):

$$\frac{\partial T/T}{\partial \theta_i} = -\frac{z_i}{\theta_i}\eta_T \tag{3}$$

$$\frac{\partial z_i}{\partial \theta_i} = -\frac{z_i}{\theta_i}\epsilon_{ii}^c \tag{4}$$

Equation (3) 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.

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 4 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 3 and 4, we can cancel (θ_i) and get an expression that links preference parameters to measurable quantities only:

$$\frac{\frac{\partial T/T}{\partial \theta_i}}{\frac{\partial z_i}{\partial \theta_i}} = \frac{\partial T/T}{\partial z_i} = \eta_T/\epsilon_{ii}^c. \tag{5}$$

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.[10] 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.

An appealing feature of equation (5) 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

[10]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.

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 5 is i specific, which means that the impact of new technologies may be different for innovations in different leisure categories. New technologies 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.

2.2 Micro Labor Supply Model

An alternative modelling approach studies individual level labor supply decisions. An advantage of this approach it is easy to introduce differences in behaviour between demographic groups, for example different behaviour over the life-cycle. Moreover, it 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 technologies that increase ν affects work hours at the intensive margin (l^*) and the retirement age at the extensive margin ($\tilde{a}$). Appendix 9.3 derives the following two comparative static results, assuming that 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. This quasi-linear utility function rules out income effects (and we discuss more general functional forms below):

$$\frac{\partial \tilde{a}}{\partial \nu} = -\frac{1}{\beta'(\tilde{a})} < 0 \tag{6}$$

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

The first result states that such technological changes lead to earlier retirements and increased exit from the labor force. Comparative static 6 shows that the retirement age declines. The new marginal retiree is thus younger, and individuals at the margin of retirement will have exited the labor force. 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. The second result shows that at the intensive margin leisure consumption increases by w^{ϵ}, with w the wage rate. The greater utility of leisure leads to a marginal reduction in work hours.

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.

3 Data

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

3.1 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.[11] 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 channels available in an area–which makes for an improvement over the binary DMA approximation of TV availability.[12]

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 FCC's signal model. This interruption 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,

[11] 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)

[12] In appendix section 10.2, we revisit the results in Gentzkow (2006) and Gentzkow and Shapiro (2008) using the new ITM data.

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.[13] 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 channels, 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 channel 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 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.[14] 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).[15] 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,

[13] 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..

[14] 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.

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

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 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."

3.2 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.[16]

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

[16]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.

includes 327,258 worker-year observations, of which people close to retirement will turn out to be particularly relevant (25,101 observations are aged 55+ in our sample window). Appendix section 11 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.

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.[17] The data reports are at annual frequency during the 1950's, however in earlier years multiple years are grouped together and multi-year summary records are available. Our baseline sample includes the annual data for 1951-1960 and two multi-year observations representing the average of 1937-1946 and 1947-1950, respectively. Our sample will span 1937 to 1960, the year our Television signal data ends.

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 11 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, 2019) for a related challenge and approach).[18] Still one may worry about at-

[17]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).

[18]Different from Kearney and Levine (2019), our results are not affected by potential migration induced changes

tenuation 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 11.3.4). 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 25,000 observations aged over 55, the sample is thin on individuals over 70 year old.[19] 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 11).

We use an additional data source to study hours worked and the intensive margin labor supply responses. In the 1950's, data on work hours was 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. Our sample includes 51 local areas and covers the period 1947-1960. The panel is thus relatively small but provides a glimpse into intensive margin effects.

4 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:

$$E_{aigt} = \gamma_{gt} + \delta_i + \beta_g \cdot TV_{at} + \pi \cdot X_{aigt} + \epsilon_{aigt}, \tag{7}$$

Here the outcome E_{aigt} is an employment indicator with value 100 if individual i of gender g in area a at time t is employed, and TV_{at} denotes the number of available TV channels 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 since these stations were typically experimental and had limited

in the composition of local labor markets since we use individual panel data and can hold individual characteristics constant.

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

broadcast hours and because many of these first launches happen before our outcome data becomes annual. Time fixed effects (γ_{tg}) absorb aggregate trends in labor supply (e.g. driven by polciies 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 main effect of interest is captured by β_g, which we allow to differ across 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 robustness checks to investigate whether parallel trends hold in this DiD set-up.

4.1 Baseline DiD Results

Using the the baseline DiD regression of equation (7), we find that an additional TV channel reduced the probability that an individual was employed by between 0.2 and 0.6 percentage points Table (1). These are relatively modestly sized effects, particularly compared to the high employment rates of 78% for men and 40% for women. While in absolute terms the estimates show larger impacts on men, baseline employment rates also higher among men. In relative terms both groups experience a similar decline in employment of roughly 0.6%. The comparable responsiveness suggests that both groups have similar underlying utility functions and share fundamental preferences. The similarity also argues against the idea that television mainly augmented home-production, as women are more engaged in home production and thus would be more exposed to televisions' effects. To interpret the magnitude of these results, it is useful to translate them into months of lost work. Below we will show that a key driver of the results are changes in retirement patterns. Suppose the effect was entirely driven by retirements then the 0.2-0.5 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).

A prominent identification concern in studies of local labor market level outcomes are changes in the composition of the local workforce. This 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 next specifications address concerns that shifting gender and family norms, expanded schooling access, more generous retirement packages, or changing life-cycle work patterns may create spurious trends. We add controls that interacted these demographic variables with linear trends and find very similar results to our baseline. We next repeat this exercise with more granular region- and state-specific time trends (Columns 4, 5 and 6). Aside from spurious trends that arise from demographic factors, these controls also capture spurious trends from unobserved factors. Column 4 introduces state specific trends and shows a qualitatively similar result but at a reduced magnitude, with significant effects for men insignificant effects on women. We unpack this reduced magnitude further since it could be a sign of two very different forces: First, the state-trends may absorb small violations of the parallel trend assumption, or second, these trends may absorb part of the dynamic treatment effects. In the former case the baseline estimates overstate the impact of TV, while in the latter case, the baseline specifications are valid and the specifications with state-trends are biased towards zero. To distinguish the two, we first fit the state specific trends using only pre-period data and thereby avoid that the trends absorbs part of the treatment effect[20] and second, we estimate trends at the more aggregated Census region level.[21] Both results are very close to the baseline DiD (Column 5 and 6) and support the baseline DiD results.[22]

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.[23] 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 11.3.4 that the magnitude of such effects is small and in a bounding exercise show that the potential impact is minor for realistic migration patterns.

[20]Specifically, we first fit state specific trends using only years without television, then extrapolate the trends to all years and residualise all variables and repeat the baseline DiD using residulized variables based on the Frisch-Waugh-Lovell theorem.

[21]Results are similar whether we use conventional linear region trends (not shown) or more flexible region-year fixed effects (Column 6).

[22]Consistent with these results, we show the dynamic pattern of the treatment effect below and confirm that labor supply adjustment is sluggish. Sharp changes at the time of the launch capture less than half of the full treatment effect (see Figure 3)

[23]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.

4.2 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.

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 parallel 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 2, 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),[24] 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 0.4 percentage point decline in employment (Column 4). This is strikingly close to our baseline and thus adds confidence in the validity of our 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

[24]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.

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 5). 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 such an analysis depends on the treatment effect's time path. If the treatment effect unfolds immediately upon impact, a narrow time window is optimal, while such a narrow window underestimates the true treatment effect if the treatment effect is slowly unfolding over time. While selecting the right window is, in principle, a thorny issue, it turns out that in our setting alternative windows yield very similar conclusion. We show results for a narrow window around the interruption period (1947-1954). The narrow window stretch our sample thin but has the advantage that they exclusively rely on years when differences in TV access are caused by the policy intervention. 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 from this experiment align closely with the baseline estimates and suggest that the results based on the overall rollout capture the causal impact of television.

4.2.1 Parallel Trends

For peace of mind, we also provide conventional pre-trend checks. The previous results and institutional features of the rollout give us reason to be optimistic that this 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 local demand conditions had no effect on the timing of television launches, giving us reason to be optimistic about parallel trends.

A recent and growing literature discusses identification strategies in two-way fixed effect settings (De Cavalcanti and Tavares, 2008; Borusyak and Jaravel, 2018; Callaway and Sant'Anna, 2020). We follow this line of work and apply the appropriate pre-trend checks developed for our setting with a continuous treatment variable and multiple event periods. There are two versions of parallel trend checks for this set-up: The first uses simple leads and lag values of our treatment and is reported in Appendix 11.3.1, the second approach uses a distributed lag model, as suggested in a series of recent work on difference-in-differences settings like ours (Fuest et al. (2015), Serrato and Zidar (2016), and Drechsler et al. (2017), and Schmidheiny and Siegloch (2019)). To implement the distributed lags regression, we uses the following first-difference transformation of equation 7:

$$\Delta E_{iagt} = \alpha_{gt} + \underbrace{\sum_{j=0}^{a} \beta_{g,j} \ \Delta \text{Channels}_{a,t-j}}_{\text{Lagged Stations}} + \underbrace{\sum_{k=1}^{a} \beta_{g,k} \ \Delta \text{Channels}_{a,t+k}}_{\text{Future Stations}} + \mathbf{\Pi} \boldsymbol{X}_{iagt} + \Delta \epsilon_{iagt} \tag{8}$$

the β_j coefficients capture the past impact of lagged stations and β_k the impact of future stations. The time pattern of a station's impact is plotted in Figure 3. The figures show that treatment and control regions evolve in parallel in the years leading up to the launch of a TV channel. The differences are close to zero and insignificant in the lead up to television launches, and after the launch of a TV station employment declines in the affected location. The effect arises at the time of treatment, and shows no sign of spurious pre-trends. The figure also shows that the treatment effect grows over time. Such slower moving effects are consistent with the notion that retirement decisions shift slowly and result in an absorbing state, where people retire once and then stay out of the labor force permanently. The cumulative effect thus increases over time, while the results do not show no signs of spurious pre-trends.

4.3 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 replicate the employment regressions in the CES data. The results show negative employment effects and broadly align with our baseline SSA results (Panel A in Table 3). 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.

4.4 Heterogeneous Effects: The Role of Retirement

We next turn to job flows and further unpack what is driving the observed results and in particular turn to Costas' hypothesis about changes in retirement behaviour (Costa, 1998).

To evaluate the retirement hypothesis more directly, Figure 4 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.[25] 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 4). These retirement effects are also substantially larger than the effects on other labor market flows. Figure 4 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 patterns also confirm the theoretical predictions derived in 2 about the impact of entertainment innovations and shows that the adjustment is largest at the retirement margin.

5 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. To evaluate the impact of the ultimate universal roll-out we need to make an extrapolation from our results since some pockets of the U.S. were still without TV at the end of our sample period in 1960. We use regression 7 to calculate the predicted employment-to-population if all areas had access to the same level of television as treated areas. This exercise suggests that television had reduced the

[25] 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.

employment to population ratio by 2% once universal coverage was achieved. Compare this effect to aggregate changes in labor supply over this time horizon: The employment to population ratio of men declined by 12 percentage points between 1950 and 1990, with a particularly steep decline of roughly 30 percentage points for individuals over the age of 55. Most of this decline happened during the first decade, which coincides with the peak television rollout. Our estimates of a 2% decline explains a sixth of the overall decline, which is a non-trivial effect but is modest relative to findings on the effect of other factors, such as the launch of the Social Security two decades earlier.[26] Importantly, female labor market trends were markedly different during this period and employment rates were growing rapidly. Our results show negative employment effects for women too, albeit smaller than for men.

This previous exercise will under-estimate the full impact of television if additional station launches after the end of our sample period had additional effects. To assess the quantitative importance of these additional later stations, we explore how the effect of stations diminishes with a growing number of available stations. It is infeasible to extend our sample to include stations outside our sample window and instead we do the opposite and test how the results change if we limit our sample to fewer launches. We estimate specifications that only use the first few station launches and compute the predicted impact of television based on these restricted samples.[27] We would expect that the predicted effects are largely unchanged if subsequent stations after the first few have limited additional effects. This is exactly what we find as shown in Figure 5. The results show that even within our sample the results are relatively unmoved by further television launches beyond the first three to five stations, suggesting that launches after the end of our sample window had minor additional effects. The figure also shows that the effect is smaller if one only uses the first station launch. In such a specification the effect is modestly negative and insignificant. The small effect of the first station may seem surprising at first but are likely driven by the extremely limited broadcast hours that were typical of pioneering stations. With the entry of competing stations hours and variety of programming expanded. In line with this, we see the employment effects grow once we take subsequent station launches into account. The employment effect stabilizes at a 2 percentage point decline in the employment-to-population ratio.

The impacts on retirement align closely with the predictions of value of leisure improvements derived above and are hard to rationalize with other potential effects of TV that do not predominantly affect retirement age cohorts. Other studies of Television show that watching Television may also affect voting, childcare, print media, teenage pregnancy, consumption of large consumer

[26] Fetter and Lockwood (2018) find that Social Security led to a decline of labor force participation among older workers of around 8.5 percentage points between 1930 and 1940.

[27] The predicted value is calculated using equation 7 and takes $E\left[E_{aigt}|TV_{at}=0\right]-E\left[E_{aigt}|TV_{at}=x\right]=\hat{\beta}_g\cdot x$, where x is the number of stations in exposed areas and grows across specifications. Note that this approach is more flexible than imposing a specific polynomial structure of effects.

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.

5.1 Leisure Innovation and Aggregate Labor Supply

The previous section focused on the impact of television, and we can leverage the results further by linking the above estimates to structural preference parameters of the representative agent model presented in Section 2.1. The chief advantage of this approach is that the model can help quantify the impact of other leisure technologies beyond television and address the question raised by Abraham and Kearney (2020) about the likely role of entertainment innovations in labor supply trends over recent decades.

Equation (5) provides the link from the reduced form estimates to structural preference parameters. We can use the above natural experiment to quantify these parameters. To use our results, we need to convert our estimates into changes in total hours worked (or equivalently total leisure time), the object modeled in the representative agent model. The 2% decline in employment implies that aggregate work hours decline by 2% and since we found no additional intensive margin effects, these estimates represent the total change in work time. The equivalent impact on total leisure time in a population of n is: $\frac{\partial T/T}{\partial \theta_i} = 0.02 \cdot 40n/31n \cdot 100 = 2.6\%$, using data on average leisure time of 31 hours (following Aguiar and Hurst (2007)) and a standard work week of 40 hours.[28] We next need to quantify $\frac{\partial z_i}{\partial \theta_i}$, the denominator of equation (5). This captures the share of leisure time devoted to home entertainment. Time use data in the pre-TV era is available in 1930 for a sample of female "homemakers." For this population the leisure share of home entertainment increased from 76% before TV to 86% after TV in 1965, implying a denominator value of equation (5) of $\frac{\partial z_i}{\partial \theta_i} = (0.86 - 0.76) = 0.1$. An alternative approach uses data from a representative cross-sectional sample in 1965. Home entertainment, excluding television, took up approximately 42% of leisure time and with television this share increases to 77%.[29] Assuming a no TV counterfactual and no other changes to leisure times, this implies a somewhat larger denominator of $\frac{\partial z_i}{\partial \theta_i} = (0.77 - 0.42) = 0.35$. Both these denominator estimates lead to the same qualitative conclu-

[28] Average hours worked are roughly 40 hours (Mcgrattan and Rogerson, 2004) and varying this assumption has little impact.

[29] In 2010 the role of television is in the same ball-park, increasing the share of home entertainment from 31% to 74%.

sion and we therefore focus on the results from the representative sample. Combining the estimates for numerator and denominator, equation (5) equals 0.074. This means that a leisure innovation that increases z_i by 10% leads to 0.7% more leisure time, approximately 14 minutes.

We can compare these estimates to existing calibrations and assess how using the quasi-experiment to identify these parameters differs from standard calibrations. The above response captures two aspects of preferences: the substitution elasticity between home entertainment and other leisure activities (ϵ_{ii}), and the substitution elasticity between leisure and consumption (η_T). 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 9.1). A large literature has estimated η_L and provides benchmark values for this parameter. For comparability, we follow Aguiar et al. (2021) and use their benchmark value of 1.1. Using these estimates in equation (5) yields $\epsilon^c_{ii} = (\eta_L \cdot \frac{L}{T}) \cdot 0.074^{-1} = (1.1 \cdot 0.78) \cdot 0.074^{-1} = 11.6$. The demand for activity i is thus relatively elastic (with an elasticity of 11.6) and people substitute easily between alternative leisure activities. An entertainment innovation for activity i thus results in relatively large changes in time spending z_i, even if total leisure time is unchanged. Note that the estimate of 11.6 is an order of magnitude larger than the calibration in Aguiar et al. (2021), which implies an elasticity of 1.62.[30] In other words, our estimates imply that leisure activities are closer substitutes than leading calibrations assume. In turn, our results imply that entertainment innovation produce smaller impacts on work time and most of the time investment will be a substitution away from other leisure activities.

The difference in results is due to the calibration choices, rather than assumptions about the utility function. While our result is based on homothetic preferences, a case with the Aguiar et al. (2021) utility assumption is presented in Appendix 9.2 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.

An alternative way to express the results is to measure how many minutes of work time are lost from one hour of television watching. For such estimates, we normalise the impact of television by the absolute amount of time investment, in minutes, instead of normalizing by the time share z_i. Such absolute time-use elasticities are arguably more tangible and easier to interpret, while capturing the same local effects as the above results. Data on time use shows the average American spent 19.3 hours watching television per week, which provides a numeraire for our estimate.[31]

[30] 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 $\epsilon^c_{ii} = \beta_i \cdot \bar{\eta} - 1 = 2.39 \cdot 1.09 - 1 = 1.62$.

[31] TV time use data is taken from the 2010 BLS series TUU10101AA01027132. TV time has been relatively stable since the 1980s and the choice of the base year thus has little impact on the results.

Using this effect as a denominator, the results imply that one hour with television reduces work by approximately by $0.02 \cdot 40/19.3 \cdot 60 = 2.5$ minutes. In other words, the vast majority of the time – over 57 minutes – of an hour spent with television crowds out other non-work activities. Leisure innovations thus appear to primarily replace alternative leisure activities.

6 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 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.

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.

Fetter, Daniel K. and Lee M. Lockwood, "Government old-age support and labor supply: Evidence from the old age assistance program," *American Economic Review*, 2018, *108* (8), 2174–2211.

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

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).

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.

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, “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.

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

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.

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.

7 Figures and Tables

Figure 1: Hours of Television Watching per Day

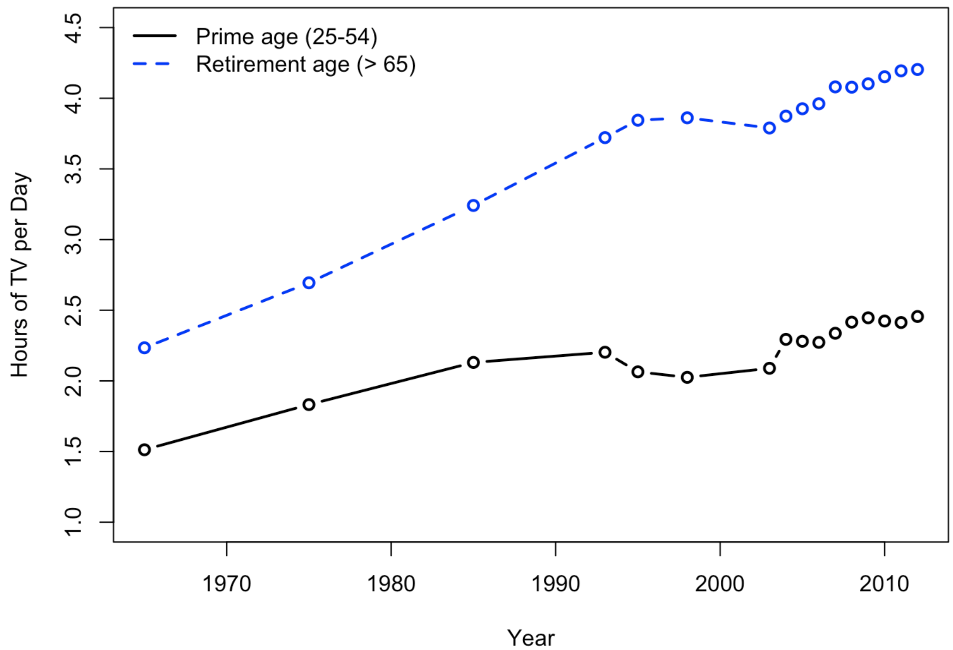

Notes: The figure shows the amount of time American's spend watching television as primary activity. Data are from the Historic American Time Use Study (AHTUS). The hours refer to "primary activity."

Figure 2: ITM-Measured Signal Strength in 1950

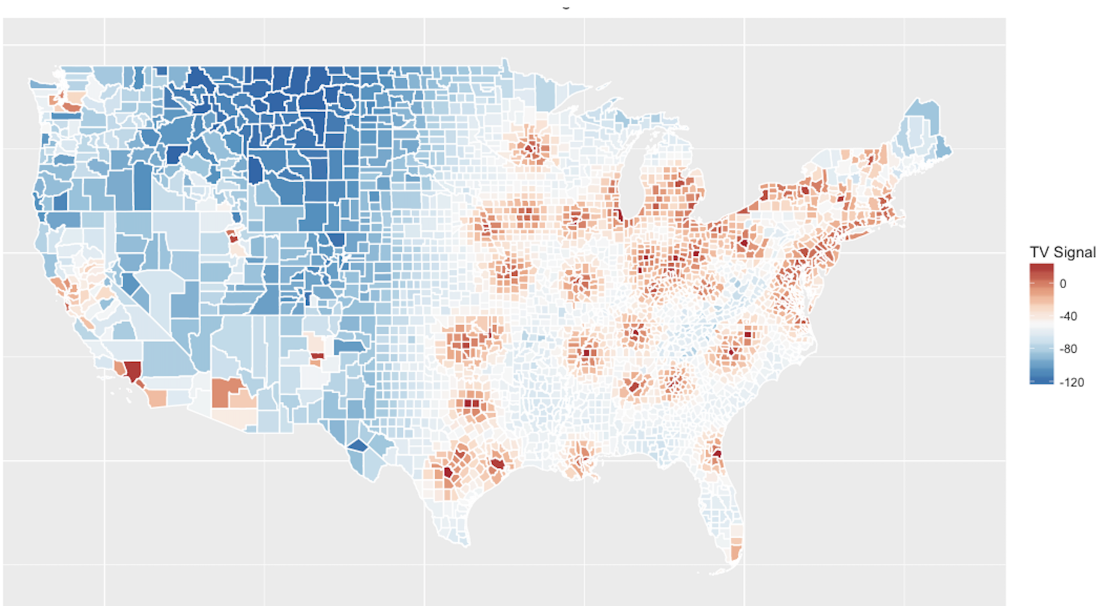

Notes: The figure shows the signal level, in decibels, of the strongest station in each county in 1950, as computed with the ITM. Broadly, counties shaded red had TV access, while counties shaded blue did not; signals whose strength was less than -50 decibels, where the map turns from red to blue, were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

Figure 3: Dynamic Effects of Station Launches

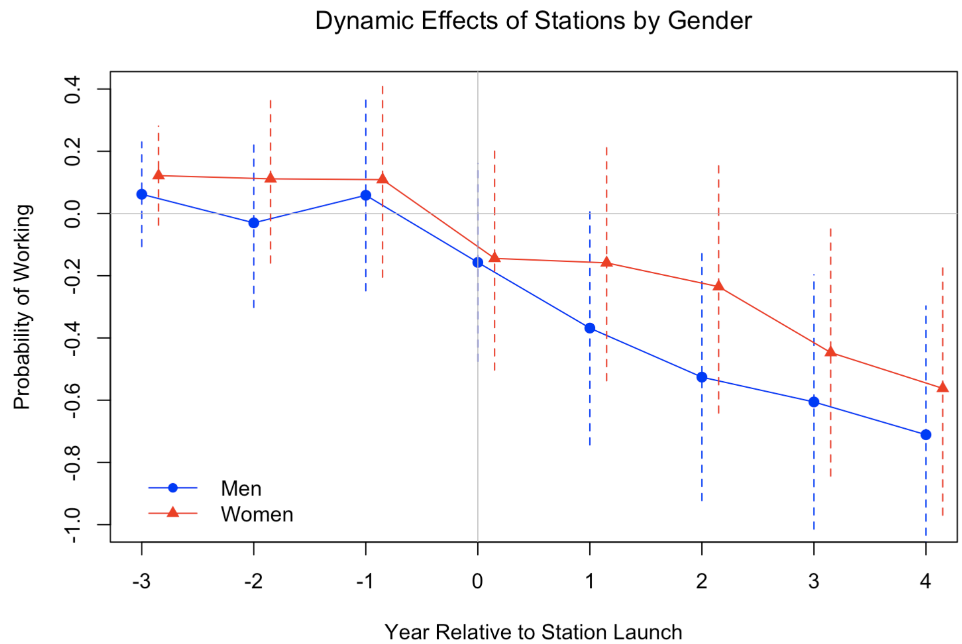

Notes: The figure shows dynamic effects of the launch of a TV station, separately for men and women. Specifically, these are estimated coefficients from individual level regression of equation 8 and 95 percent confidence intervals, following the approach in Serrato and Zidar (2016) and related recent work with a continuous treatment variable. See text for further details.

Figure 4: Effects of TV on Entry, Exit, and Retirement

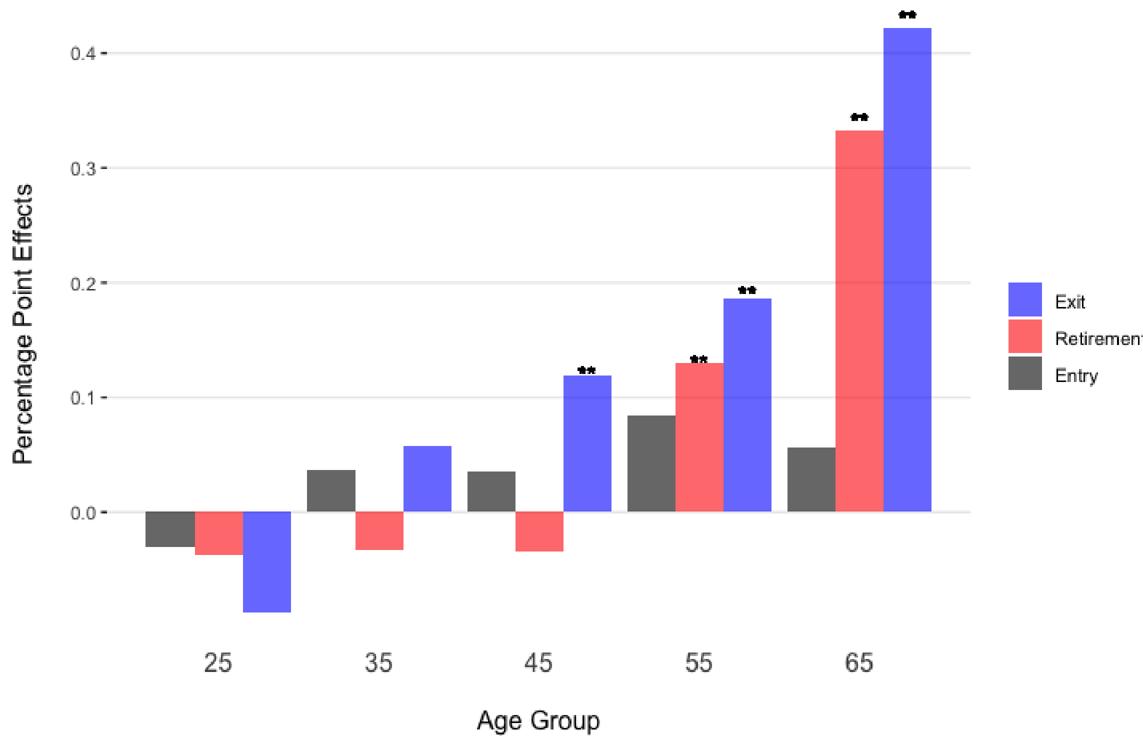

Notes: The figure shows the impact of television on job transitions. Effects on employment entry are shown in black, on exits in blue and on retirement in red. Retirement is an indicator with value hundred in the year a worker exits permanently from employment (proxied by the absence of a work observation 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 286,698 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 1 percent level.

Figure 5: Steady State Effect of TV Accounting for Additional Stations

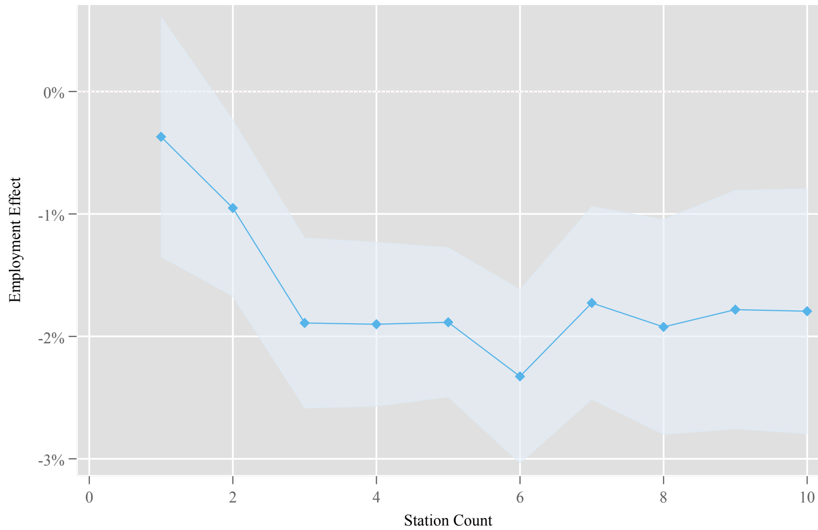

Notes: The figure shows the predicted impact of television based on the predicted effects ($\hat{y}$) from regression 7 for the average treated area. We calculate the predicted impact using the full sample and for cases were we only use data on the the first, the first two, etc. station launches in each area. We estimate 7 for each of these samples and show the number of launches in the sample on the x-axis and the predicted effect on the y-axis. A stable predicted effect implies that taking additional station launches into account has little impact on the predicted long-run impact of television. The shaded are 90% confidence bands.

Table 1: Individual-level Effects of TV on Employment

	(1)	(2)	(3)	(4)	(5)	(6)
$\mathbb{1}$(Men) × Stations	-0.574***	-0.585***	-0.600***	-0.315**	-0.590***	-0.681***
	(0.131)	(0.132)	(0.130)	(0.140)	(0.133)	(0.187)
$\mathbb{1}$(Women) × Stations	-0.246**	-0.246**	-0.261**	0.0281	-0.250**	-0.222*
	(0.111)	(0.110)	(0.112)	(0.122)	(0.110)	(0.123)
Observations	325,130	325,130	325,130	325,130	325,130	325,130
R-squared	0.678	0.679	0.680	0.680	0.688	0.679
Year × Gender FE	Yes	Yes	Yes	Yes	Yes	Yes
Person FE	Yes	Yes	Yes	Yes	Yes	Yes
Age FE	No	Yes	Yes	Yes	Yes	Yes
Trends	No	No	Demographics	State	State*	Region × Year
Mean DV Men	78.29	78.29	78.29	78.29	78.29	78.29
Mean DV Women	38.28	38.28	38.28	38.28	38.28	38.28

Notes: The table shows individual level regressions of an employment dummy with value 100 for an employed worker on the number of TV stations available in the local area. 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. Demographic trends allow for different time trends for high-school graduates, race (white, black, other), marital status and 5 year age bins. "State*" controls for state specific trends, where the state trends are estimated using pre-TV data only (see text for details). Regions are census regions. Television is measured at the MSA level. Standard errors are clustered at the same level and span 134 clusters. Source: SSA-CPS employment records and Television Factbooks $***p < 0.01, **p < 0.05, *p < 0.1$

Table 2: 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	
Blocked stations	0.120*	0.107				
	(0.0703)	(0.0707)				
Stations	-0.321***	-0.335***	-0.386***	-0.391***	-0.447***	-0.419***
	(0.0978)	(0.0979)	(0.0987)	(0.0993)	(0.112)	(0.112)
Observations	317,016	317,016	257,856	257,856	99,644	99,644
R-squared	0.680	0.681	0.680	0.680	0.775	0.775
Demographic 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, 1947-1954). The estimates use the baseline specification in column 3 of Table 1. See Table notes for additional details $***p < 0.01, **p < 0.05, *p < 0.1$

Table 3: 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

8 Appendix B: Representative Agent Framework

Here we solve the utility optimisation problem in 2 and derive predictions about leisure demand. We can solve the utility maximisation problem in two-steps.[32] 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:

$$V(T,\boldsymbol{\theta}) = \max_{\alpha_i} \quad v(\alpha_1, ..., \alpha_n) \quad \text{s.t. } T = \sum_1^n \alpha_i/\theta_i \tag{9}$$

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 $\theta_i^- 1 \equiv \tilde{\theta}_i$, 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 $\frac{c}{w} + P \cdot V = 1$.[33] 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:

$$V(1, P(\tilde{\theta}_1/T, ..., \tilde{\theta}_n/T)) \tag{10}$$

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 $V(1, P(\tilde{\theta}_1/T, ..., \tilde{\theta}_n/T))$.[34]

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 10: $\frac{\partial V}{\partial \tilde{\theta}_i} = \frac{1}{T}\frac{\partial V}{\partial P}\frac{\partial P}{\partial \tilde{\theta}_i}$ 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:

$$\alpha_i = -\frac{\frac{\partial V}{\partial \tilde{\theta}_i}}{\frac{\partial V}{\partial T}} = \frac{T \cdot \partial V/\partial P \cdot \partial P/\partial \tilde{\theta}_i}{\kappa} = \Pi_i T \tag{11}$$

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

[33] 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.

[34] We can use a similar monotonic transformation to obtain a homogenous function representation of the preferences represented by the homothetic function $v(\boldsymbol{a})$.

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 11 with $\frac{\tilde{\theta}_i}{T}$ and using $t_i = \alpha_i \cdot \tilde{\theta}$, we find z_i:

$$z_i = \frac{\alpha_i \cdot \tilde{\theta}}{T} = \Pi_i\tilde{\theta}_i = \frac{\partial P}{\partial\tilde{\theta}_i}\frac{\tilde{\theta}_i}{P} \tag{12}$$

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).

9 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:

$$\frac{\partial T}{\partial\tilde{\theta}_j}\frac{\tilde{\theta}_j}{T} = \frac{\partial T}{\partial P}\frac{P}{T}\frac{\partial P}{\partial\tilde{\theta}_j}\frac{\tilde{\theta}_j}{P} = z_j\eta_T \tag{13}$$

where the last equality uses 12 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 3 in the main text.

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 11 yields a solution for ϵ^c_{ii}:

$$\epsilon^c_{ii} = 1 + \frac{\partial \Pi_i}{\partial \tilde{\theta}_i} \frac{\tilde{\theta}_i}{\Pi_i} \tag{14}$$

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 12 and using 14 and 12 we find:

$$\frac{\partial z_i}{\partial \tilde{\theta}_i} = \Pi_i + \frac{\partial \Pi_i}{\partial \tilde{\theta}_i} \tilde{\theta}_i = \Pi_i \epsilon^c_{ii} = z_i/\tilde{\theta}_i \epsilon^c_{ii} \tag{15}$$

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 4 in the main text by multiplying both sides of 15 with $\frac{1}{\theta_j^2}$ and using the fact that $\frac{\partial z_i}{\partial \tilde{\theta}_j} = -\frac{\partial z_i}{\partial \theta_j} \theta_j^2$.

9.1 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

$$\eta_T = \frac{\partial T}{\partial P} \frac{P}{T} = -\frac{\partial c}{\partial P} \frac{1}{w} \frac{P}{T}. \tag{16}$$

The Slutsky equation implies:

$$\frac{\partial c}{\partial P} = \frac{\partial c}{\partial P}|_C + c\frac{\partial c}{\partial m} = \frac{\partial V}{\partial 1/w}|_C + c\frac{\partial c}{\partial m}. \tag{17}$$

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:

$$\frac{\partial V}{\partial 1/w} = -\frac{w^2}{P}\frac{\partial L}{\partial w} \tag{18}$$

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 17 by assuming $\frac{\partial L}{\partial m} = 0$, in line with a sizable labor supply literature that finds modest income effects.[35] Using this in 17 and combining the results with 16 and 18, 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.

9.2 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:

$$v(\alpha_1, ..., \alpha_i) = \sum_i \frac{\alpha_i^{1-1/\eta_i}}{1 - 1/\eta_i} \tag{19}$$

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

$$\mathcal{L} = v(\alpha_1, ..., \alpha_i) - \mu\left[\sum \frac{\alpha_i}{\theta_i} - T\right] \tag{20}$$

[35]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 $\frac{\partial V}{\partial 1/w} = \frac{\partial V}{\partial 1/w}|_C$.

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

$$\begin{aligned} FOC_i: \quad & ln(t_i) = (\eta_i - 1)\cdot ln(\theta_i) - \eta_i \cdot ln(\mu) \\ FOC_\mu: \quad & T = \textstyle\sum_i t_i \end{aligned} \tag{21}$$

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:

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

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:

$$\frac{\partial ln(\mu)}{\partial ln(T)} = -\frac{1}{\bar{\eta}} \tag{23}$$

We can obtain the compensated demand elasticity of t_i by differentiating FOC_i and using 22:

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

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:

$$\frac{\partial ln(t_i)}{\partial ln(1/\theta_i)}|_u \equiv \epsilon^u_{ii} = -(\eta_i - 1)(1 - \beta_i \cdot z_i) - \beta_i \frac{ln(T)}{ln(1/\theta_i)} \tag{25}$$

We can use the envelope theorem to obtain two further useful results. Evaluating 20 at the optimal α_i^* and differentiating with respect to T and θ_i respectively yields the following results:

$$\begin{aligned} V_T &= \mu \\ V_T \cdot t_i &= V_{\theta_i} \cdot \theta_i \end{aligned} \tag{26}$$

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 2 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:

$$\begin{aligned} \tfrac{\partial ln(u_v)}{\partial ln(T)} + \tfrac{\partial ln(V_T)}{\partial ln(T)} &= \tfrac{\partial ln(\Omega)}{\partial ln(T)} \\ \tfrac{\partial ln(u_v)}{\partial ln(\theta_i)} + \tfrac{\partial ln(V_T)}{\partial ln(\theta_i)} &= \tfrac{\partial ln(\Omega)}{\partial ln(\theta_i)} \end{aligned} \tag{27}$$

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

$$\begin{aligned} \tfrac{\partial ln(u_v)}{\partial ln(T)} &= \tfrac{T}{u_v}\left[u_{vv}V_T + u_{vc}\tfrac{\partial c}{\partial T}\right] \\ \tfrac{\partial ln(u_v)}{\partial ln(\theta_i)} &= \tfrac{\theta_i}{u_v}\left[u_{vv}V_{\theta_i} + u_{vc}\tfrac{\partial c}{\partial \theta_i}\right] \end{aligned} \tag{28}$$

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:

$$\begin{aligned} 0 &= u_{cc}\tfrac{\partial c}{\partial T} + u_{vc}V_T \\ 0 &= u_{cc}\tfrac{\partial c}{\partial \theta_i} + u_{vc}V_{\theta_i} \end{aligned} \tag{29}$$

Combining 29, 28 and 27 yields:

$$\left[\frac{\partial ln(V_T)}{\partial ln(T)} - \frac{\partial ln(\Omega)}{\partial ln(T)}\right]\frac{V_{\theta_i}\theta_i}{V_T T} + \frac{\partial ln(V_T)}{\partial ln(\theta_i)} = \frac{\partial ln(\Omega)}{\partial ln(\theta_i)} \tag{30}$$

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 26 into equation 30 yields:

$$\left[\frac{\partial ln(\mu)}{\partial ln(T)} - \frac{1}{e}\right] z_i + \frac{\partial ln(\mu)}{\partial ln(\theta_i)} = \frac{\partial ln(\Omega)}{\partial ln(\theta_i)} \tag{31}$$

Combining this with 22 and 23 yields:

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

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:

$$\frac{\partial ln(T)}{\partial ln(\theta_i)} = \frac{\partial ln(T)}{\partial ln(\Omega)} \frac{\partial ln(\Omega)}{\partial ln(\theta_i)} = e \cdot \frac{\beta_i \cdot e - 1}{e} z_i = (\beta_i \cdot e - 1) z_i \tag{33}$$

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 5 in the main text. The change in z_i is:

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

we can now derive the equivalent to the main result in 5 for the non-homothetic preference case. The relative change in T compared to z_i is:

$$\frac{\partial T/T}{\partial z_i} = \frac{\frac{\partial ln(T)}{\partial ln(\theta_i)}}{\frac{\partial ln(z_i)}{\partial ln(\theta_i)} \cdot z_i} = \frac{1}{\left(\frac{\partial ln(t_i)}{\partial ln(\theta_i)} / \frac{\partial ln(T)}{\partial ln(\theta_i)} - 1\right) z_i} \tag{35}$$

using 33 and 25, we can obtain the equivalent expression to 5, for the utility function used by Aguiar et al. (2021):

$$\frac{\partial T/T}{\partial z_i} = \frac{1}{\left(\frac{-\epsilon^c_{ii}}{(\beta_i \cdot e - 1) z_i} - (1 - \beta_i)\right) z_i} \tag{36}$$

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 (5).[36] Otherwise, in the presence of non-homothetic

[36] 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.

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_{ii}^c:

$$\epsilon_{ii}^c = [1 + \psi \cdot (1 - \beta_i) z_i] \cdot \frac{\beta_i \cdot e - 1}{\psi} \tag{37}$$

For small values of ψ — as the case in our empirical application — the first term is ≈ 1 and 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_{ii}^c 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_{ii}^c by close to 2.4 times.

9.3 Individual Level Labor Supply Model

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:

$$max\ U(c, \xi(a)l) \tag{38}$$

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

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 $\tilde{a}$:

$$\xi(\tilde{a}) = \frac{x(1+\epsilon)}{w^{1+\epsilon}} \tag{39}$$

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 39 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.

The simplicity of the results hinges on the functional form assumption, but some of these predictions hold more broadly. Intensive margin results on l^* are sensitive to parametric assumptions. If individuals have a strong income effect, the direction of the change could go the other way and the impact of entertainment technologies at the intensive margin in the general model is thus ambiguous. This highlights one of the problems with testing intensive margin effects of entertainment technologies. Studies typically assume that income effects are small or absent to arrive at unambiguous predictions about l^*. Our extensive margin predictions, by contrast, are not sensitive to the functional form assumptions. These results are one of the few predictions of the general labor supply framework that hold independently of the parametric assumptions about the utility function.

10 Appendix C: Measuring TV Access

10.1 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.

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

Proximal cities confound the DMA approximation of TV access. For example, panel (A) of figure 11 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.[37]

[37]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

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.

10.1.2 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 11 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 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

signal strength. County coverage exceeding 50 percent for a station means that over 50 percent of households in the county watched that channel. Our measurement of signal reach will not rely on these maps.

that are associated with TV ownership.

10.1.3 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[38]. 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.[39]

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 13, 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.[40] 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 11 examples show with Kansas City and Minneapolis-St. Paul, and the fringe plays a key role in estimating TV's effects. Gentzkow (2006) exploits 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

[38]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.

[39]Table 8 lists the first 40 pairs and shows the distance between towers. Note also that in 1948 the FCC froze applications for new broadcast licenses in part because it realized it had allowed stations to be too close together.

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

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 3.1 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.

10.2 TV Data Validation Exercise

As referenced in the introduction, much of our knowledge on the effects of TV relies on the DMA approximation. Among the many papers using the DMA approach are Baker and George (2010) on household debt, Campante and Hojman (2013) on political polarization, Thomas (2019) on smoking, Kim (2020) on consumer culture and spending, and Angelucci et al. (2021) on media competition and news consumption. The original DMA papers are Gentzkow (2006) and Gentzkow and Shapiro (2008) on how TV impacted voter turnout and children's test scores, respectively. Here we replicate the main results of these two papers using the ITM, and we find that the estimated effects are about twice as large with the new data.[41]

Gentzkow (2006) studies how the 1950's TV rollout affected voter turnout. The direction of the effect is a priori ambiguous–it could be that TV broadened news viewership and therefore stimulated political engagement, or, alternatively, that TV crowded out news consumption with entertainment programming, which in turn dampened political knowledge and interest. Gentzkow finds robust evidence for the the latter case, using the following baseline difference-in-differences specification:

$$Y_{it} = \alpha_i + \delta_{rt} + \gamma TV_{it} + \beta X_{it} + \epsilon_{it} \tag{40}$$

Here the outcome Y_{it} is voter turnout in county i and year t, and controls include county effects α_i, region-year effects δ_{rt}, as well as flexible time trends interacted with county characteristics in X_{it}. The explanatory variable of interest TV_{it} is the number of years that county i has had TV access in year t, so the coefficient γ captures the effect of an additional year of TV access on voter turnout.

Row 1 of table 9 reports the main results from the paper. Columm 2, the fully-controlled, preferred specification shows that an additional year of TV availability led to 0.136 percentage point decline in voter turnout, an effect size that "explains half of the total off-year decline in turnout since the 1950's. The effect on presidential-year turnout is smaller—accounting for

[41] We are grateful to Matthew Gentzkow for his correspondence and generous assistance with code and data.

roughly a quarter of the total decline— and is not significantly different from zero" (p. 933). (Note that the effects in row 1 are much larger for the column 4 mid-term elections than the column 3 presidential elections.) Rows two and three show results using the ITM rather than the DMA's to measure TV access, with both a -40 and -50 decibel threshold for access. The effects are upwards of 2-3 times larger, which is consistent with a reduction in attenuation bias arising from measurement error.

We find similar results in the context of a study on TV and education. Gentzkow and Shapiro (2008) investigate how TV influenced children's test scores, providing a rigorous test of longstanding worries that TV could "rot children's brains" using data from the 1965 Coleman Report. This paper uses a two-stage least squares approach, instrumenting for TV ownership in a household with the availability of a TV signal; the idea is that TV ownership and viewership may well have been endogenous choices, but that conditional on a set of controls, access to a TV signal was idiosyncratic. The central results are based on the following first- and second-stage regressions:

$$y_{gc} = \beta TV_{gc} + \phi_g W_c + \delta_c + \gamma_g + \epsilon_{gc} \tag{41}$$

$$TV_{gc} = \beta_g^0 ADOPT_c + \phi_g^0 W_c + \delta_c^0 + \gamma_g^0 + \epsilon_{gc}^0 \tag{42}$$

The main outcome y_{gc} in equation 41 is average test scores for students in grade g and location c, which is regressed on the number of years of potential preschool television exposure for those students, TV_{gc}, and additional controls. Gentzkow and Shapiro (2008) instrument for TV_{gc} in equation 42 with a variable $ADOPT_c$ for the time at which location c adopted TV broadcasts, as measured using the DMA approximation.

Table 10 reports the main findings from the paper, as well as the first-stage F-statistic from equation 42. Contrary to common narratives about the harmful influence of TV, the row 1 results show that, if anything, TV exposure during childhood *increased* test scores. Many of the effects are imprecise, but they are positive, and for reading scores, the coefficient is statistically significant, "consistent with a variety of existing evidence suggesting that children can learn language-based skills from television" (p. 300). In rows 2 and 3, we estimate the same two-stage least squares specification using the ITM to measure TV access. Note first that first-stage F-statistic is larger, meaning there is a stronger association between TV signal availability and TV ownership using the ITM. We take this as validation that the ITM is more accurately measuring signal reach than the DMA's. The effects on test scores in columns 2-5 are larger and more precise as well, with the exception of general knowledge scores.

Taken together, these replication exercises suggest that future researchers studying the effects of TV should use the ITM measurements of access. The DMA approach appears to produce substantial underestimates of TV's influence. We aim to make the ITM data available for both further revisions of existing results and future original work.

11 Appendix B: Empirical Appendix

11.1 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.[42] 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.[43]

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

[42]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.

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

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.

11.2 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 4. 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.

11.3 Robustness Checks

11.3.1 Leads and Lags

A popular method to check for pre-trends is to include leads and lags of the treatment in the event study designs and analyze changes in labor supply in the lead up to an event. The intiuition is that effects should arise after television launch events and not before. We implement this through a dynamic DiD which replicates DiD 7 and additionally allows for leads and lags of the treatment:

$$E_{a,i,t} = \gamma_t + \delta_i + \sum_{j=-4}^{3} \beta_{t+j} \cdot TV_{a,t+j} + \pi \cdot X_{a,i,t} + \epsilon_{a,i,t},$$

these leads and lags capture the evolution of the treatment effect in the 4 years before and after the launch of a new TV channel. We exclude the two multi-year bins from this analysis, since the leads and lags for multi-year windows is not well defined. Conventional event studies have to omit one lead or lag regressor, because these regressors are otherwise co-linear with the year dummies. In our case, we have one more degree of freedom because the television treatment varies in intensity. More than one television station is launched in some years, which breaks the perfect co-linearity of leads and lags and time FE. We could therefore estimate coefficients for all lead and lag periods but since we are mainly interested in trends, we follow the standard event-study design and normalise the effects in period t-1 to zero.[44] This eases the interpretation of the results, as coefficients then capture the deviation in employment relative to the t-1 period. Table 5 shows that treatment and control regions evolve in parallel in the years leading up to the launch of a TV channel. And we see a sharp change after the launch of a TV station. The clear change at the time of treatment indicates that the difference-in-difference specification is capturing the effects of TV and we can rule out that differences in trends are driving our results. The following columns control for alternative aggregate and regional trends and find similar results.

11.3.2 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.[45] Specifically, we target population aggregates in an MSA, as well as their education and age demographics.

[44]The effect in t-1 is typically positive around 0.1, reflecting that places with multiple simultaneous launches have higher rates of employment. For display purposes, we purge the impact of this level effect and subtract this value from all coefficients.

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

Table 6 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 are are slightly larger for men and smaller for women but the differences are not statistically significant. Similarly, if we use weights for the steady-state estimates, we again find consistent results (see Figure 10). The impact of television increases with the first few stations and then steadies out. The point estimate is a 1.3% decline in the employment to population ratio, close to the baseline 1.8% estimate and the difference is insignificant. The weighted and unweighted results thus show qualitatively and quantitatively similar results.

11.3.3 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 7, 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 age differences and again highlights that the effect is much bigger among workers near retirement. Column 3 and 4 look at effects by schooling and marital status. The effect on both groups is similar to the baseline estimates.

11.3.4 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 7 reports treatment effects for this non-complier group in column 1. Using $\sigma = 0.95$ and $NCTE = -0.301$ in the ATT formula yields an ATT of -0.397, very close to the ITT estimate of -0.392.

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 -0.408, 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=-0.431, 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.

12 Appendix Figures and Tables

Figure 6: Marginal Retiree

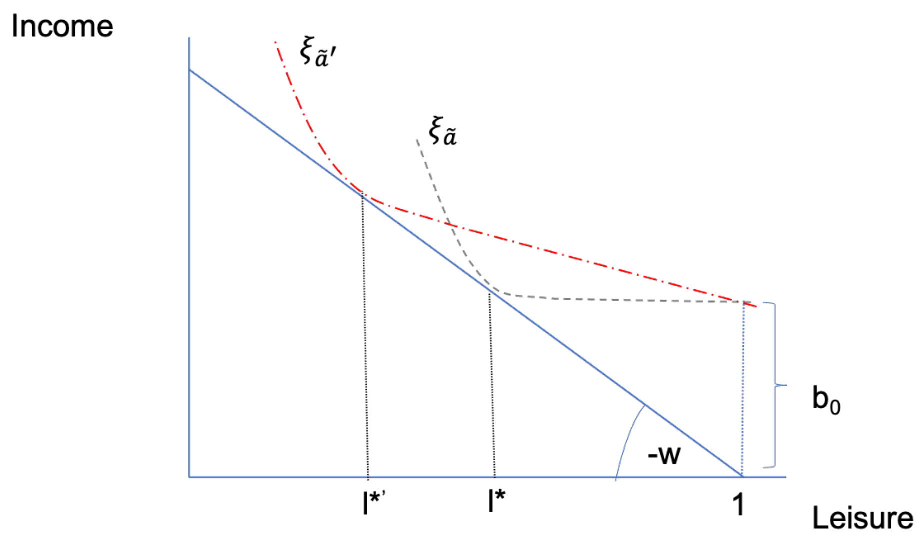

Notes: The figure shows the indifference curve of the marginal retiree, a person who is just indifferent between working and not. The age of the marginal retiree is indicated by $\tilde{a}$. The dashed line is a case with low β_0 and the dash-dot line is a case with higher β_0.

Figure 7: Retirement Rates

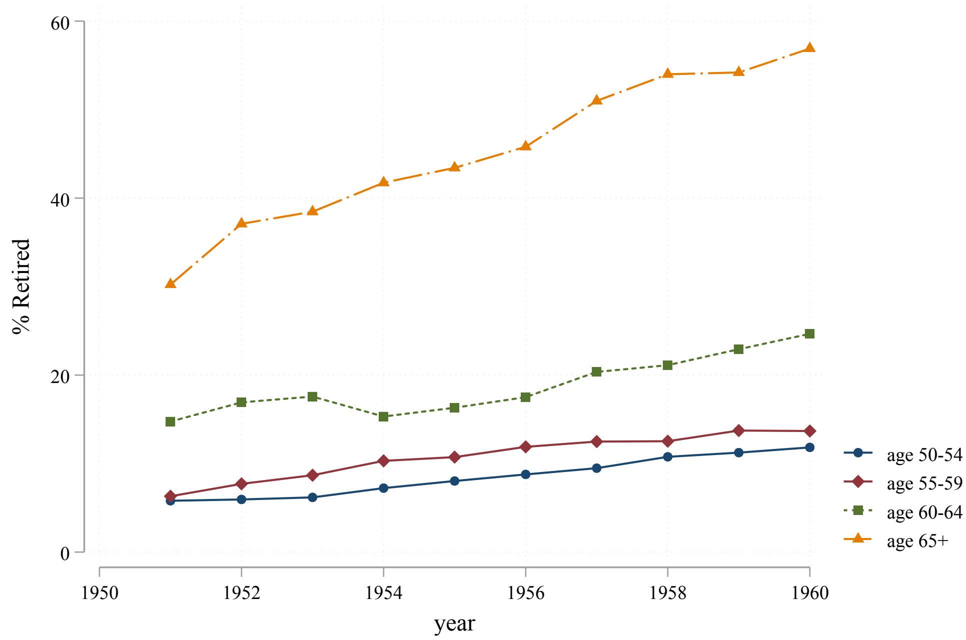

Notes: The figure shows retirement rates among older workers during the 1950's. Retirement is defined as no observed employment in the Social Security records until the end of our sample (1978). Source: linked SSA-CPS data.

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

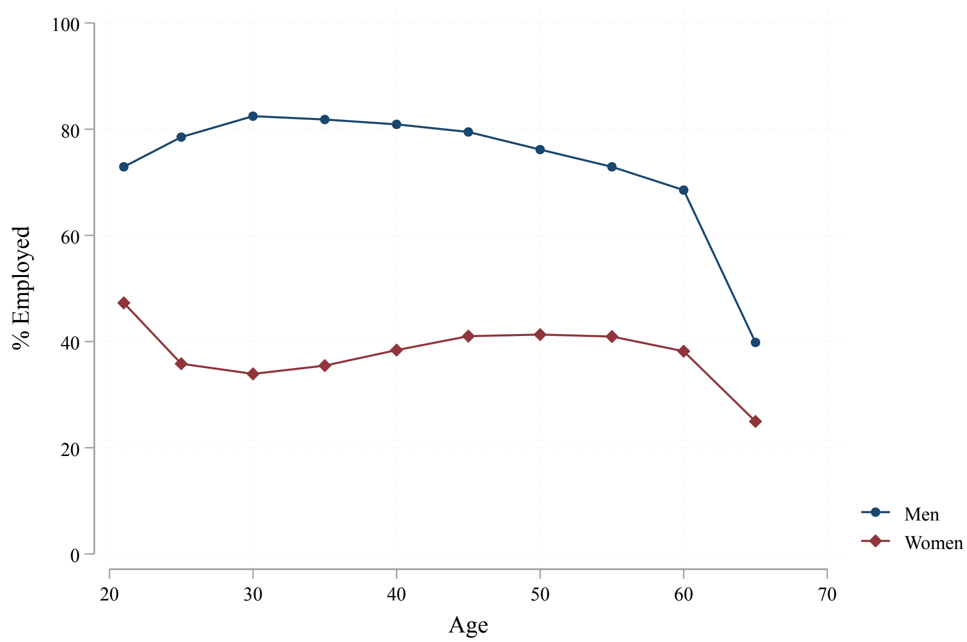

Notes: The figure shows employment rates by age and gender. Each dot shows the average for a five year age window, averaging employment rates over the full sample period. The first and last bins respectively show averages for the age groups 21-24 years and 65+. Source: linked SSA-CPS data.

Figure 9: Number of Stations Available Over Time

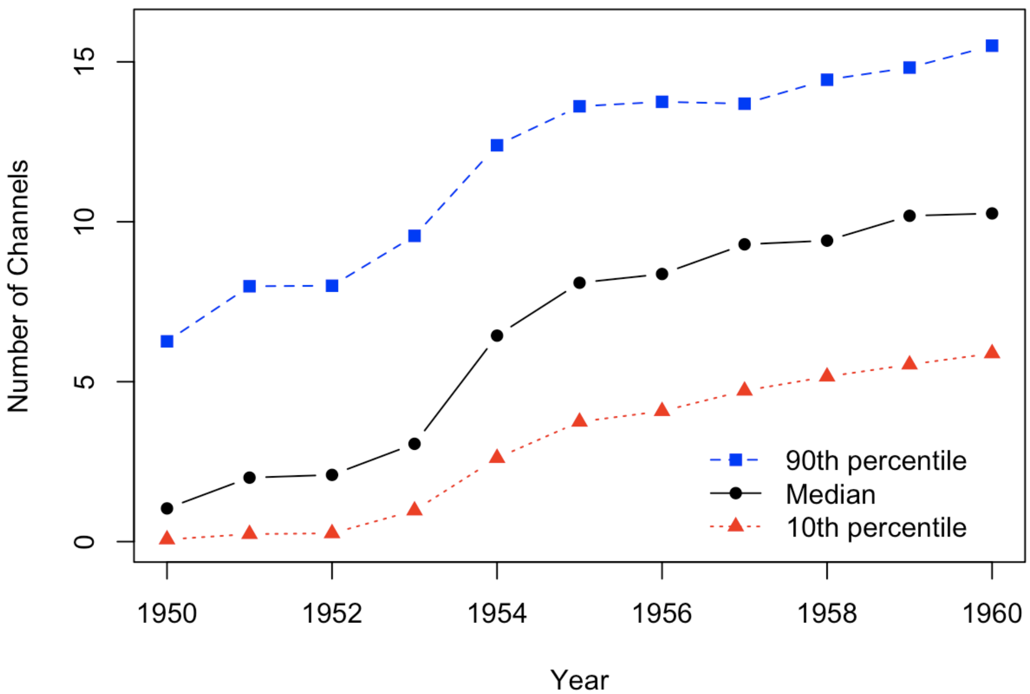

Notes: The figure shows the number of television stations in the U.S. between 1950 and 1960. It shows this for a median person, as well as at the 90th and 10th percentile of the distribution.

Figure 10: Steady State Effect of TV Accounting for Additional Stations - Weighted Sample

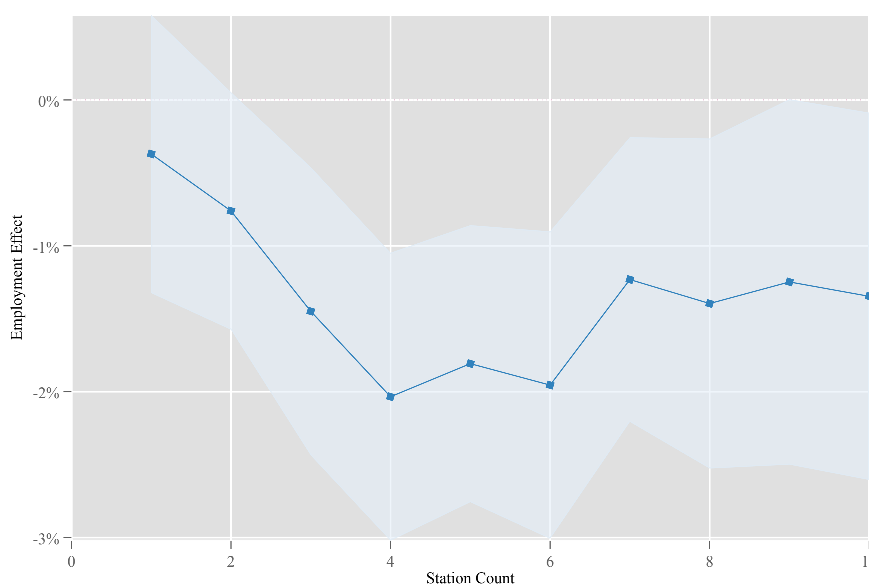

Notes: The figure replicates Figure 5 while using sample weights. Weights are constructed to make the sample representative of local population demographics at the annual level.

Figure 11: Coverage Maps and Designated Market Areas

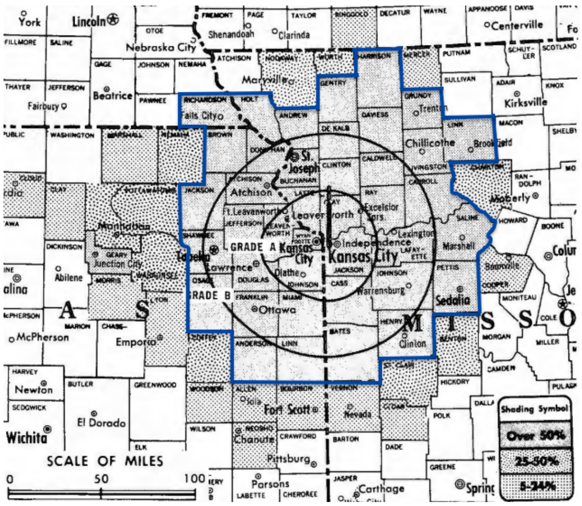

(A) Kansas City coverage map ring (in blue)

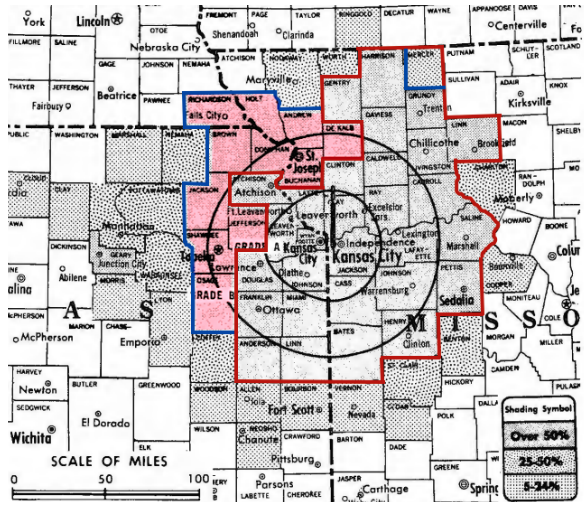

(B) Kansas City DMA ring (in red)

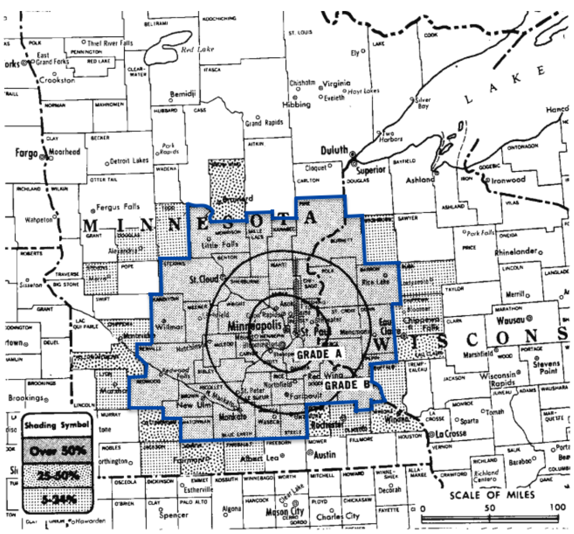

(C) Minneapolis coverage map ring (in blue)

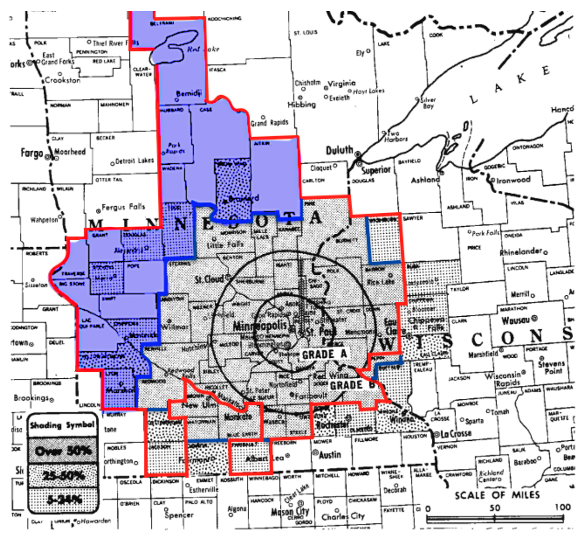

(D) Minneapolis DMA ring (in red)

Figure 12: TV Purchases Patterns

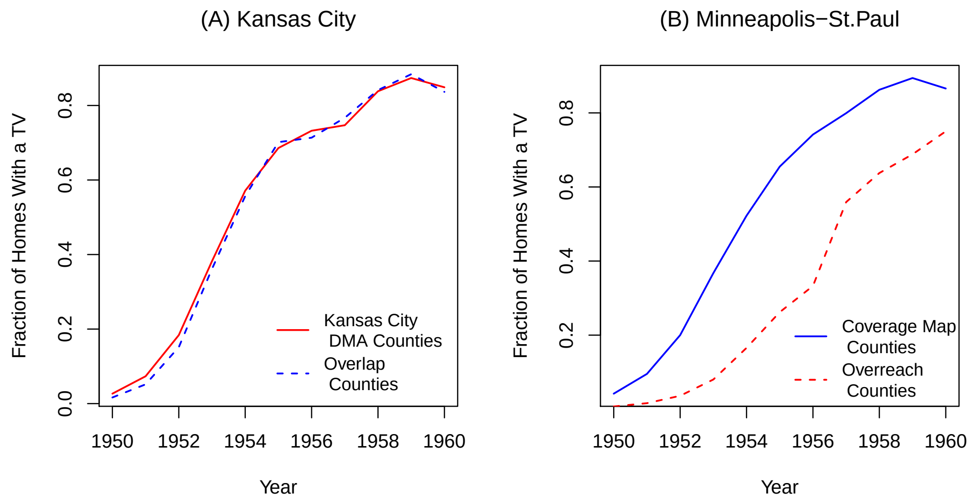

Notes: Panel (A) shows average TV ownership around Kansas City for counties in the groups indicated in the legend. "Overlap Counties" refers to those highlighted in Figure . In Panel (B), for Minneapolis-St. Paul, "Coverage Map Counties" refers to those ringed in Figure 4, whose coverage exceeds 50 percent according to *TV Factbook* coverage maps. "Overreach Counties" refers to those highlighted in Figure 5, which fall inside the Minneapolis-St. Paul DMA but outside the *TV Factbook* broadcast range.

Figure 13: Broadcast Technology Improvements

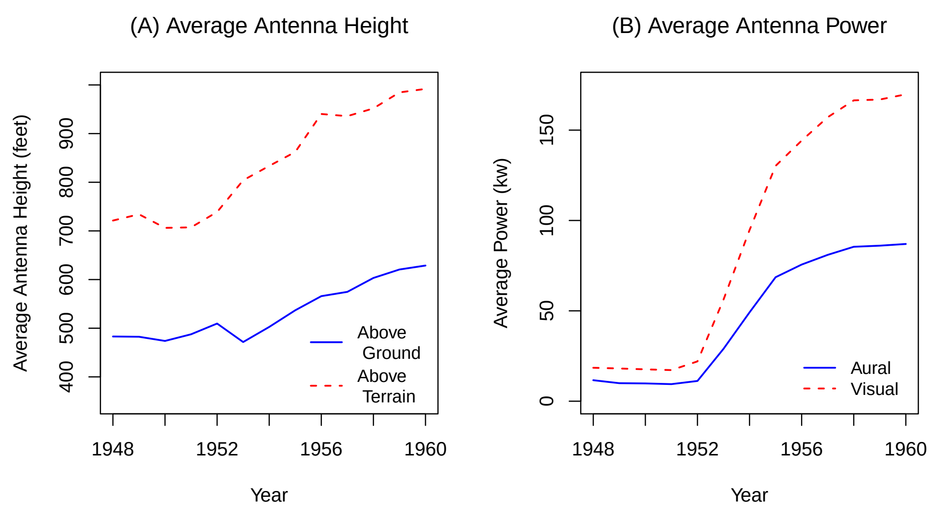

Notes: The figure shows the increases in broadcast tower height and power over time. Data are digitized from the *TV Factbook*, as discussed in the main text.

Table 4: Summary Statistics

	Observation	Average	s.d.	Min	Max	Men	Women
Employed	325,130	54.54	49.79	0	100	78.29	38.28
Quarters worked	325,042	1.909	1.877	0	4	2.886	1.239
TV channels	325,130	6.904	4.697	0	16.65	6.910	6.899
Years of schooling	325,130	11.80	3.419	1	19	11.69	11.87
High school graduate	325,130	0.541	0.498	0	1	0.508	0.563
Year of Birth	325,130	1916	11.24	1881	1938	1916	1917
Ever married	325,130	0.950	0.217	0	1	0.947	0.953
Female	325,130	0.594	0.491	0	1	0	1
Minority	325,130	0.0883	0.284	0	1	0.0922	0.0855
Recent move	321,196	0.0521	0.222	0	1	0.0526	0.0518

Notes: The table reports summary statistics for the SSA-CPS sample. Employment and age information is based on SSA records and spans the years 1937-1960. The data is annual from 1951 to 1960 and includes 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 channels is computed using records from digitized Television Factbooks in an ITM signal propagation model.

Table 7: Heterogeneous Effects of TV on Employment by Demographic Groups

	(1)	(2)	(3)	(4)
Stations	-0.392***	-0.345***	-0.432***	-0.466***
	(0.0977)	(0.0966)	(0.101)	(0.144)
Stations $\times$ $\mathbb{1}$(Mobile person)	0.0878			
	(0.141)			
Stations $\times$ $\mathbb{1}$(Age 60+)		-0.584***		
		(0.133)		
Stations $\times$ $\mathbb{1}$(High school dropout)			0.0874*	
			(0.0500)	
Stations $\times$ $\mathbb{1}$(Married)				0.0849
				(0.121)
Observations	322,139	326,089	326,089	326,089
R-squared	0.680	0.680	0.680	0.680

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 moved MSA between 1975 and 1976. Robust standard errors in parentheses. $***p < 0.01, **p < 0.05, *p < 0.1$

Table 5: TV Effects on Employment – Leads and Lags

	(1)	(2)	(3)	(4)	(5)	(6)
t-4	0.005	0.001	-0.002	0.221	-0.043	-0.001
	(0.122)	(0.123)	(0.125)	(0.155)	(0.146)	(0.141)
t-3	-0.103	-0.109	-0.108	-0.096	-0.222	-0.096
	(0.119)	(0.119)	(0.120)	(0.145)	(0.134)	(0.132)
t-2	0.003	0.008	0.001	0.037	-0.219	0.064
	(0.0959)	(0.0966)	(0.0970)	(0.115)	(0.0945)	(0.106)
t-1	0	0	0	0	0	0
t	-0.273	-0.274	-0.268	-0.256	-0.371	-0.238
	(0.105)	(0.104)	(0.105)	(0.103)	(0.113)	(0.112)
t+1	-0.247	-0.240	-0.238	-0.177	-0.277	-0.177
	(0.107)	(0.107)	(0.107)	(0.117)	(0.124)	(0.125)
t+2	-0.256	-0.260	-0.259	-0.192	-0.359	-0.156
	(0.116)	(0.116)	(0.117)	(0.118)	(0.139)	(0.118)
t+3	-0.265	-0.253	-0.246	-0.057	-0.380	-0.269
	(0.129)	(0.128)	(0.128)	(0.168)	(0.135)	(0.135)
Observations	161,483	161,483	161,483	161,483	161,483	161,483
R-squared	0.782	0.782	0.782	0.782	0.783	0.782
Cluster	134	134	134	134	134	134
Year $\times$ Sex FE	Yes	Yes	Yes	Yes	Yes	Yes
Person FE	Yes	Yes	Yes	Yes	Yes	Yes
Age FE	No	Yes	Yes	Yes	Yes	Yes
Trends	None	None	Demographics	State	State*	Region $\times$ Year

Notes: The Table shows the timing of television effects by reporting coefficients on the leads and lags of the television variable. Period $t-1$ is normalised to 0 to illustrate changes in the effect around the time of television launches. See Table 1 for variable definitions and additional specification details.

Table 6: Individual-level Effects of TV on Employment - Weighted Sample

	(1)	(2)	(3)	(4)	(5)	(6)
$\mathbb{1}$(Men) $\times$ Stations	-0.749***	-0.752***	-0.728***	-0.444**	-0.756***	-0.829***
	(0.154)	(0.155)	(0.154)	(0.173)	(0.155)	(0.206)
$\mathbb{1}$(women) $\times$ Stations	-0.110	-0.105	-0.126	0.169	-0.109	-0.0819
	(0.170)	(0.169)	(0.170)	(0.172)	(0.169)	(0.172)
Sum of Weights (thsd.)	1,953,494	1,953,494	1,953,494	1,953,494	1,953,494	1,953,494
R-squared	0.678	0.679	0.680	0.680	0.687	0.680
Year $\times$ Sex FE	Yes	Yes	Yes	Yes	Yes	Yes
Person FE	Yes	Yes	Yes	Yes	Yes	Yes
Age FE	No	Yes	Yes	Yes	Yes	Yes
Trends	No	No	Demographics	State	State*	Region $\times$ Year
Mean DV Men	78.5	78.5	78.5	78.5	78.5	78.5
Mean DV Women	38.5	38.5	38.5	38.5	38.5	38.5

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

Figure 14: Revisiting TV's Effects on Voter Turnout (Gentzkow, 2006)

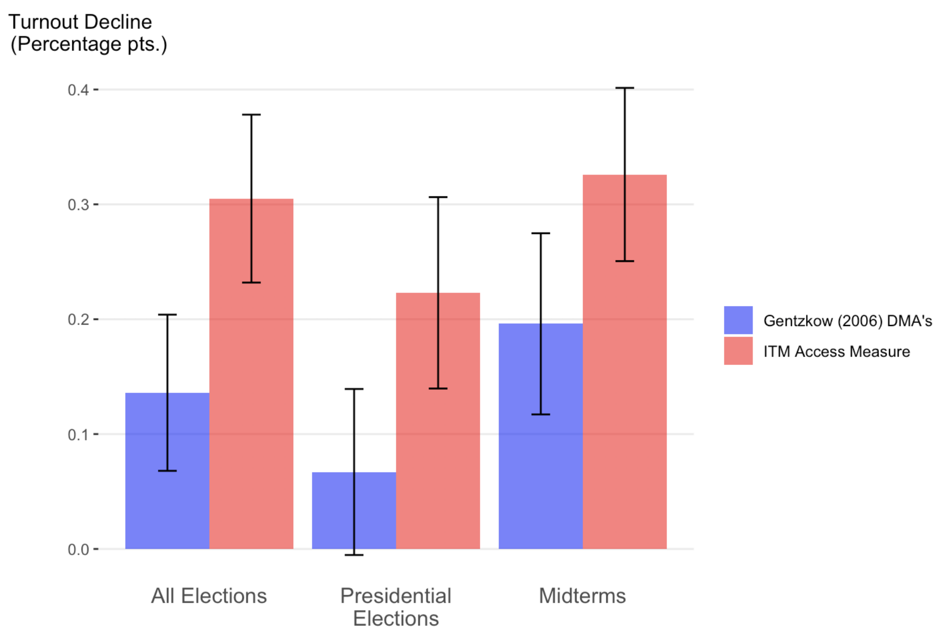

Table 8: Proximal Market Areas

DMA 1	DMA 2	Miles Apart	Years Apart
Pittsburgh (PA) [1949]	Steubenville (OH) [1954]	32.79	5
Washington (DC) [1946]	Harrisburg (PA) [1949]	35.86	3
Harrisonburg (VA) [1954]	Charlottesville (VA) [1960]	36.04	6
Harrisburg (PA) [1949]	Johnstown (PA) [1950]	42.47	1
Cleveland (OH) [1948]	Youngstown (OH) [1953]	42.53	5
Grand Rapids (MI) [1949]	Lansing (MI) [1950]	45.46	1
Binghamton (NY) [1950]	Elmira (NY) [1953]	45.67	3
Syracuse (NY) [1949]	Utica (NY) [1950]	46.36	1
Kansas City (MO) [1950]	St. Joseph (MO) [1954]	48.35	4
Cincinnati (OH) [1948]	Dayton (OH) [1949]	48.48	1
Lake Charles (LA) [1954]	Beaumont (TX) [1955]	49.55	1
Youngstown (OH) [1953]	Steubenville (OH) [1954]	50.28	1
Columbus (OH) [1949]	Zanesville (OH) [1953]	52.28	4
Binghamton (NY) [1950]	Wilkes Barre (PA) [1953]	52.39	3
Zanesville (OH) [1953]	Parkersburg (WV) [1954]	52.44	1
Cleveland (OH) [1948]	Steubenville (OH) [1954]	52.49	6
Detroit (MI) [1947]	Toledo (OH) [1948]	53.08	1
San Francisco (CA) [1949]	Sacremento (CA) [1954]	54.15	5
Baton Rouge (LA) [1953]	Lafayette (LA) [1955]	54.94	2
Pittsburgh (PA) [1949]	Youngstown (OH) [1953]	57.01	4
Hartford (CT) [1948]	Springfield (MA) [1953]	57.39	5
Nashville (TN) [1951]	Bowling Green (KY) [1960]	58.19	9
Grand Rapids (MI) [1949]	South Bend (IN) [1953]	58.36	4
Indianapolis (IN) [1949]	Lafayette (IN) [1953]	58.74	4
Lima (OH) [1953]	Ft. Wayne (IN) [1954]	58.86	1
Kansas City (MO) [1950]	Topeka (KS) [1954]	59.70	4
South Bend (IN) [1953]	Ft. Wayne (IN) [1954]	60.10	1
Birmingham (AL) [1949]	Montgomery (AL) [1953]	60.13	4
Memphis (TN) [1949]	Jonesboro (AR) [1960]	60.48	11
Jacksonville (FL) [1950]	Gainesville (FL) [1960]	61.83	10
Roanoke (VA) [1953]	Charlottesville (VA) [1960]	62.10	7
Denver (CO) [1952]	Colorado Springs (CO) [1953]	63.65	1
Rochester (MN) [1953]	La Crosse (WI) [1954]	63.69	1
Richmond (VA) [1948]	Norkfolk (VA) [1950]	63.88	2
Washington (DC) [1946]	Baltimore (MD) [1948]	63.95	2
Champaign (IL) [1953]	Terre Haute (IN) [1954]	64.67	1
Syracuse (NY) [1949]	Watertown (NY) [1955]	65.18	6

Notes: In brackets is the year in which a broadcast began in each DMA. Some DMAs are abbreviated for brevity. For example, the Birmingham (AL) - Anniston (AL) - Tuscaloosa (AL) DMA is listed just as Birmingham (AL).

Table 9: Revisiting TV's Effects on Voter Turnout (Gentzkow, 2006)

	All Elections	All Elections	Presidential	Non-presidential
DMA	-0.416	-0.136	-0.067	-0.196
	(0.0486)	(0.0412)	(0.0438)	(0.0478)
ITM_{40}	-0.468	-0.254	-0.171	-0.278
	(0.0450)	(0.0421)	(0.0481)	(0.0438)
ITM_{50}	-0.513	-0.305	-0.223	-0.326
	(0.0479)	(0.0443)	(0.0505)	(0.0457)
Full controls		X	X	X

Notes: The table replicates the Gentzkow (2006) results on TV's influence on voter turnout, with both the original DMA approximation and the new ITM data. ITM_{40} and ITM_{50} refer to measurements of TV access using -40 and -50 decibel cutoffs for access, respectively. Column 2 is the preferred specification in the paper, which shows effects on the order of 2-3 larger using the ITM. Column 3 shows results for the sub-sample of presidential election years, column 4 for off-presidential mid-term elections. See figure 14 for a plot of the DMA and ITM_{50} coefficients and 90 percent confidence intervals.

Table 10: Revisiting TV's Effects on Children's Test Scores (Gentzkow and Shapiro, 2008)

	First Stage F Stat.	Average Score	Verbal	Reading	General Knowledge
DMA	16.58	0.0225	0.0294	0.0557	0.0672
		(0.0279)	(0.0289)	(0.0302)	(0.0410)
ITM_{40}	36.69	0.0385	0.0511	0.0598	0.0384
		(0.0200)	(0.0214)	(0.0247)	(0.0310)
ITM_{50}	23.87	0.0374	0.0485	0.0604	0.0338
		(0.0231)	(0.0238)	(0.0276)	(0.0376)
Full controls		X	X	X	X

Notes: The table revisits the Gentzkow and Shapiro (2008) findings on how TV affected children's test scores. As before, ITM_{40} and ITM_{50} refer to measurements of TV access using -40 and -50 decibel cutoffs for access, respectively, while DMA refers to the DMA approximation to TV braodcast reach. These are two-stage least squares estimates, where TV ownership is instrumented with TV access; the first-stage F-statistic shows how strongly the reported measures of TV access predict TV ownership. See figure 15 for a plot of the DMA and ITM_{50} coefficients and 90 percent confidence intervals.

Figure 15: Revisiting TV's Effects on Children's Test Scores (Gentzkow and Shapiro, 2008)