

ONLINE APPENDIX FOR

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

Authors: George Fenton & Felix Koenig

A Representative Agent Framework

We provide a representative agent framework that closely mirrors prior work in the area and builds on the model [Aguiar et al. \(2021\)](#) use to study the effect of videogames.¹ The model helps understand how aggregate labor supply in the economy (i.e. total work hours) responds to innovations that enhance the value of leisure time.

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

$$u(c, v(\alpha_1, \dots, \alpha_n)) \quad (\text{A.1})$$

where $v(\alpha_1, \dots, \alpha_n)$ 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 B.2](#), finding small effects on the predictions we study. We therefore focus on the simpler homogenous utility case first.

The utility maximisation problem of this agent is:

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

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.² 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.³ 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 .

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 A.1](#) for derivations):

$$\frac{\partial T/T}{\partial \theta_i} = -\frac{z_i}{\theta_i} \eta_T \quad (\text{A.3})$$

$$\frac{\partial z_i}{\partial \theta_i} = -\frac{z_i}{\theta_i} \epsilon_{ii}^c \quad (\text{A.4})$$

¹Similarly, a large literature in macro-economics uses this approach (for recent examples on labor supply, see [Rachel \(2020\)](#); [Boppart and Krusell \(2020\)](#))

²Note 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.

³From $L + T = 1$ it follows that the elasticity of leisure demand (η_T) is proportional to the labor supply elasticity (η_L) with proportionality factor $\frac{-T}{L}$: $\eta_L = \eta_T \frac{-T}{L}$.

Equation (A.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 . It might look as if we could avoid this issue by bringing θ_i to the left side of the equation. However, this implies that we can no longer estimate the left hand side. The LHS would be the percent change in leisure time in response to a given percent change in the value of leisure. Since we don't know the latter, we cannot estimate this object.

We propose an alternative approach that is independent of technology units and expresses individual responses as "time-use elasticities." These elasticities normalise the reduced form change in work time by a numeraire that captures the magnitude of the technology shock and in our case the numeraire is z_i – the time share of leisure activities i . z_i provides a revealed preference metric for the quality improvement that results from a new technology since people will spend more time with a more appealing entertainment activity. The impact of θ_i on z_i is shown in equation A.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 A.3 and A.4, we can cancel (θ_i) and get an expression that links preference parameters to measurable quantities only:

$$\frac{\frac{\partial T/T}{\partial \theta_i}}{\frac{\partial z_i}{\partial \theta_i}} = \frac{\partial T/T}{\partial z_i} = \eta_T / \epsilon_{ii}^c. \quad (\text{A.5})$$

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 .⁴ 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 (A.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 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 A.5 is i specific, which means that the impact of new technology may be different for innovations in different leisure categories. New technology in home entertainment could have a different effect than a similar innovation in sport activities, and there is a separate sufficient statistic for each group of leisure activities.

⁴Quantifying 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.

A.1 Derivation

We can solve the utility maximisation problem in A.2 in two-steps.⁵ 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:

$$V(T, \boldsymbol{\theta}) = \max_{\alpha_i} v(\alpha_1, \dots, \alpha_n) \quad \text{s.t. } T = \sum_1^n \alpha_i / \theta_i \quad (\text{A.6})$$

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

The second step of the utility maximisation problem maximises the upper nest of the utility function. Using the indirect utility function $V(T, P) = v(\alpha_1^*, \dots, \alpha_n^*)$ and defining $P = f(\boldsymbol{\theta})$ as the price index for a unit of V , the budget constraint for the upper level problem is $\frac{c}{w} + P \cdot V = 1$.⁶ 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:

$$V(1, P(\tilde{\theta}_1/T, \dots, \tilde{\theta}_n/T)) \quad (\text{A.7})$$

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 $V(1, P(\tilde{\theta}_1/T, \dots, \tilde{\theta}_n/T))(\frac{1}{T})^\beta$ or $V(1, P(\tilde{\theta}_1/T, \dots, \tilde{\theta}_n/T))$.⁷

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 A.7: $\frac{\partial V}{\partial \theta_i} = \frac{1}{T} \frac{\partial V}{\partial P} \frac{\partial P}{\partial \theta_i}$ 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:

$$\alpha_i = -\frac{\frac{\partial V}{\partial \theta_i}}{\frac{\partial V}{\partial T}} = \frac{T \cdot \frac{\partial V}{\partial P} \cdot \frac{\partial P}{\partial \tilde{\theta}_i}}{\kappa} = \Pi_i T \quad (\text{A.8})$$

where the last step combines all the terms that are independent of T into $\Pi_i \equiv \frac{\frac{\partial P}{\partial \tilde{\theta}_i}}{\sum_{i=1}^n \frac{\partial P}{\partial \tilde{\theta}_i} \tilde{\theta}_i} = \frac{\partial \ln(P)}{\partial \tilde{\theta}_i}$, 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 A.8 with $\frac{\tilde{\theta}_i}{T}$ and using $t_i = \alpha_i \cdot \tilde{\theta}_i$, we find z_i :

$$z_i = \frac{\alpha_i \cdot \tilde{\theta}_i}{T} = \Pi_i \tilde{\theta}_i = \frac{\partial P}{\partial \tilde{\theta}_i} \frac{\tilde{\theta}_i}{P} \quad (\text{A.9})$$

⁵The proof of this claim is available upon request.

⁶This 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.

⁷We can use a similar monotonic transformation to obtain a homogenous function representation of the preferences represented by the homothetic function $v(\boldsymbol{\alpha})$.

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).

B 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 :

$$\frac{\partial T}{\partial \tilde{\theta}_j} \frac{\tilde{\theta}_j}{T} = \frac{\partial T}{\partial P} \frac{P}{T} \frac{\partial P}{\partial \tilde{\theta}_j} \frac{\tilde{\theta}_j}{P} = z_j \eta_T \quad (\text{B.1})$$

where the last equality uses [A.9](#) 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 $\frac{T}{\tilde{\theta}_j}$ yields [A.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 $\frac{\partial \ln(t_i)}{\partial \ln(\tilde{\theta}_i)} = 1 + \frac{\partial \ln(\alpha_i)}{\partial \ln(\tilde{\theta}_i)} \equiv \epsilon_{ii}^c$ and denote this compensated time use elasticity for activity i by ϵ_{ii}^c . Combining this result with the derivative of the log of [A.8](#) yields a solution for ϵ_{ii}^c :

$$\epsilon_{ii}^c = 1 + \frac{\partial \Pi_i}{\partial \tilde{\theta}_i} \frac{\tilde{\theta}_i}{\Pi_i} \quad (\text{B.2})$$

Π_i embodies the marginal value of leisure and the result shows that the elasticity of Π_i is equal to $\epsilon_{ii}^c - 1$. 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 [A.9](#) and using [B.2](#) and [A.9](#) we find:

$$\frac{\partial z_i}{\partial \tilde{\theta}_i} = \Pi_i + \frac{\partial \Pi_i}{\partial \tilde{\theta}_i} \tilde{\theta}_i = \Pi_i \epsilon_{ii}^c = z_i / \tilde{\theta}_i \epsilon_{ii}^c \quad (\text{B.3})$$

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

We can now derive equation A.4 in the main text by multiplying both sides of B.3 with $\frac{1}{\theta_j^2}$ and using the fact that $\frac{\partial z_i}{\partial \theta_j} = -\frac{\partial z_i}{\partial \theta_j} \theta_j^2$.

B.1 The Elasticity of Leisure and Labor

If we are further interested in separately identifying η_T and ϵ_{ii}^c , 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

$$\eta_T = \frac{\partial T}{\partial P} \frac{P}{T} = -\frac{\partial c}{\partial P} \frac{1}{w} \frac{P}{T}. \quad (\text{B.4})$$

The Slutsky equation implies:

$$\frac{\partial c}{\partial P} = \frac{\partial c}{\partial P} \Big|_C + c \frac{\partial c}{\partial m} = \frac{\partial V}{\partial 1/w} \Big|_C + c \frac{\partial c}{\partial m}. \quad (\text{B.5})$$

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:

$$\frac{\partial V}{\partial 1/w} = -\frac{w^2}{P} \frac{\partial L}{\partial w} \quad (\text{B.6})$$

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 B.5 by assuming $\frac{\partial L}{\partial m} = 0$, in line with a sizable labor supply literature that finds modest income effects.⁸ Using this in B.5 and combining the results with B.4 and B.6, we get:

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

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.

⁸note 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 $\frac{\partial V}{\partial 1/w} = \frac{\partial V}{\partial 1/w} \Big|_C$.

B.2 Non-Homothetic Preferences

We now explore the impact of relaxing the homotheticity assumption on $v(a_1, \dots, a_i)$. We follow the approach in [Aguilar et al. \(2021\)](#) and assuming the following functional form:

$$v(\alpha_1, \dots, \alpha_i) = \sum_i \frac{\alpha_i^{1-1/\eta_i}}{1 - 1/\eta_i} \quad (\text{B.7})$$

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

$$\mathcal{L} = v(\alpha_1, \dots, \alpha_i) - \mu \left[\sum \frac{\alpha_i}{\theta_i} - T \right] \quad (\text{B.8})$$

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

$$\begin{aligned} FOC_i : \quad \ln(t_i) &= (\eta_i - 1) \cdot \ln(\theta_i) - \eta_i \cdot \ln(\mu) \\ FOC_\mu : \quad T &= \sum_i t_i \end{aligned} \quad (\text{B.9})$$

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

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

$$\frac{\partial \ln(\mu)}{\partial \ln(\theta_i)} = \frac{\eta_i - 1}{\bar{\eta}} z_i \quad (\text{B.10})$$

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 FOC s and differentiating with respect to T yields:

$$\frac{\partial \ln(\mu)}{\partial \ln(T)} = -\frac{1}{\bar{\eta}} \quad (\text{B.11})$$

We can obtain the compensated demand elasticity of t_i by differentiating FOC_i and using [B.10](#):

$$\frac{\partial \ln(t_i)}{\partial \ln(1/\theta_i)} \Big|_c \equiv \epsilon_{ii}^c = -(\eta_i - 1)(1 - \beta_i \cdot z_i) \quad (\text{B.12})$$

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:

$$\frac{\partial \ln(t_i)}{\partial \ln(1/\theta_i)} \Big|_u \equiv \epsilon_{ii}^u = -(\eta_i - 1)(1 - \beta_i \cdot z_i) - \beta_i \frac{\ln(T)}{\ln(1/\theta_i)} \quad (\text{B.13})$$

We can use the envelope theorem to obtain two further useful results. Evaluating [B.8](#) at the

optimal α_i^* and differentiating with respect to T and θ_i respectively yields the following results:

$$\begin{aligned} V_T &= \mu \\ V_T \cdot t_i &= V_{\theta_i} \cdot \theta_i \end{aligned} \quad (\text{B.14})$$

where $V = v(\alpha_1^*, \dots, \alpha_i^*)$ 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 A.2 at $V(T, P)$ and taking first order conditions:

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

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 behaviour of demand, log differentiate FOC_T with respect to T and to θ_i :

$$\begin{aligned} \frac{\partial \ln(u_v)}{\partial \ln(T)} + \frac{\partial \ln(V_T)}{\partial \ln(T)} &= \frac{\partial \ln(\Omega)}{\partial \ln(T)} \\ \frac{\partial \ln(u_v)}{\partial \ln(\theta_i)} + \frac{\partial \ln(V_T)}{\partial \ln(\theta_i)} &= \frac{\partial \ln(\Omega)}{\partial \ln(\theta_i)} \end{aligned} \quad (\text{B.15})$$

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

$$\begin{aligned} \frac{\partial \ln(u_v)}{\partial \ln(T)} &= \frac{T}{u_v} \left[u_{vv} V_T + u_{vc} \frac{\partial c}{\partial T} \right] \\ \frac{\partial \ln(u_v)}{\partial \ln(\theta_i)} &= \frac{\theta_i}{u_v} \left[u_{vv} V_{\theta_i} + u_{vc} \frac{\partial c}{\partial \theta_i} \right] \end{aligned} \quad (\text{B.16})$$

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

$$\begin{aligned} 0 &= u_{cc} \frac{\partial c}{\partial T} + u_{vc} V_T \\ 0 &= u_{cc} \frac{\partial c}{\partial \theta_i} + u_{vc} V_{\theta_i} \end{aligned} \quad (\text{B.17})$$

Combining B.17, B.16 and B.15 yields:

$$\left[\frac{\partial \ln(V_T)}{\partial \ln(T)} - \frac{\partial \ln(\Omega)}{\partial \ln(T)} \right] \frac{V_{\theta_i} \theta_i}{V_T T} + \frac{\partial \ln(V_T)}{\partial \ln(\theta_i)} = \frac{\partial \ln(\Omega)}{\partial \ln(\theta_i)} \quad (\text{B.18})$$

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 B.14 into equation B.18 yields:

$$\left[\frac{\partial \ln(\mu)}{\partial \ln(T)} - \frac{1}{e} \right] z_i + \frac{\partial \ln(\mu)}{\partial \ln(\theta_i)} = \frac{\partial \ln(\Omega)}{\partial \ln(\theta_i)} \quad (\text{B.19})$$

Combining this with B.10 and B.11 yields:

$$\frac{\partial \ln(\Omega)}{\partial \ln(\theta_i)} = \frac{\beta_i \cdot e - 1}{e} z_i \quad (\text{B.20})$$

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:

$$\frac{\partial \ln(T)}{\partial \ln(\theta_i)} = \frac{\partial \ln(T)}{\partial \ln(\Omega)} \frac{\partial \ln(\Omega)}{\partial \ln(\theta_i)} = e \cdot \frac{\beta_i \cdot e - 1}{e} z_i = (\beta_i \cdot e - 1) z_i \quad (\text{B.21})$$

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

$$\frac{\partial \ln(z_i)}{\partial \ln(\theta_i)} = \frac{\partial \ln(t_i)}{\partial \ln(\theta_i)} - \frac{\partial \ln(T)}{\partial \ln(\theta_i)} \quad (\text{B.22})$$

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

$$\frac{\partial T/T}{\partial z_i} = \frac{\frac{\partial \ln(T)}{\partial \ln(\theta_i)}}{\frac{\partial \ln(z_i)}{\partial \ln(\theta_i)} \cdot z_i} = \frac{1}{\left(\frac{\partial \ln(t_i)}{\partial \ln(\theta_i)} / \frac{\partial \ln(T)}{\partial \ln(\theta_i)} - 1 \right) z_i} \quad (\text{B.23})$$

using B.21 and B.13, we can obtain the equivalent expression to A.5, for the utility function used by Aguiar et al. (2021):

$$\frac{\partial T/T}{\partial z_i} = \frac{1}{\left(\frac{-\epsilon_{ii}^c}{(\beta_i \cdot e - 1) z_i} - (1 - \beta_i) \right) z_i} \quad (\text{B.24})$$

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 (A.5).⁹ Otherwise, in the presence of non-homothetic preferences, there is an additional income effect that affects the elasticity of z_i . The magnitude of this effect depends on β_i ; an entertainment innovation produces larger changes in z_i among leisure luxuries ($\beta_i > 1$) than among leisure necessities ($\beta_i < 1$).

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

$$\epsilon_{ii}^c = [1 + \psi \cdot (1 - \beta_i) z_i] \cdot \frac{\beta_i \cdot e - 1}{\psi} \quad (\text{B.25})$$

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

⁹Both e and η_T are closely linked to the “Frisch” elasticity of labor supply. However the link is derived under different assumptions. e assumes that u_c is unaffected by θ_i , e.g. with additive separable preferences, while η_T directly maps into the “Frisch elasticity” if there are no income effects. We therefore need the assumption about the values of e and η_T to make the two models coincide.

the main impact of β_i is to scale the denominator of the final expression up or down. For typical parameter estimates, this effect is modest relative to the gaps we identified in the text. [Aguilar 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 [Aguilar et al. \(2021\)](#) ($\beta_i = 2.4$) increases the estimate of ϵ_{ii}^c by close to 2.4 times.

B.3 Individual Level Labor Supply Model

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

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

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

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

$$\max U(c, \xi(a)l) \tag{B.26}$$

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

Denote the value of leisure of a person is just indifferent between working and no by $\xi(\tilde{a})$. [Figure E.1](#) 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 $\tilde{a}$:

$$\xi(\tilde{a}) = \frac{x(1 + \epsilon)}{w^{1+\epsilon}} \quad (\text{B.27})$$

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 B.27 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 E.1 shows the intuition behind this result. The rising value of β_0 pivots the indifference curve upward and makes it steeper. This implies that the new marginal retiree has $\xi(\tilde{a}') < \xi(\tilde{a})$, and hence $\tilde{a}' < \tilde{a}$. The new marginal retiree is thus younger, and individuals with age between $\tilde{a}'$ and $\tilde{a}$ will have exited the labor force.

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

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

C Measuring TV Access

C.1 ITM Model Details

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

over counties.

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

C.2 Measurement Error in the DMA Data

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

C.2.1 An Example of DMA Underestimation (A type II error)

Proximal cities confound the DMA approximation of TV access. For example, panel (A) of figure [E.5](#) 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.¹⁰

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

¹⁰For two reasons, the *Factbook* maps ought to be taken only as suggestive regarding true 1950's signal reach. The first is that these maps were not published until the 1960's, and tower technology—power, height, etc.—improved substantially over time. The second is that the shading in the maps reflects surveys of viewership, not measures of signal strength. County coverage exceeding 50 percent for a station means that over 50 percent of households in the county watched that station. Our measurement of signal reach will not rely on these maps.

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.

C.2.2 An Example of DMA Overestimation (A type I error)

Today's DMAs sometimes extend further from city centers than historical signals did. Panel (C) of figure E.5 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 that are associated with TV ownership.

C.2.3 Causes and Prevalence of Measurement Error

This section moves beyond examples to the causes of measurement error and evidence on the prevalence of those causes. To start with underestimation, the two conditions under which the signal overlap problem arises are: Neighboring DMA towers (1) are close enough for signals to overlap and (2) started broadcasts in different years¹¹. The closer the towers and the further apart the initial broadcast years, the larger the potential measurement error. To find possible areas of overlap, we ranked pairs of DMAs by their distance apart. There are 166 unique pairs of DMAs whose towers are less than 100 miles apart (a typical broadcast radius) with broadcasts beginning in different years. Among them are the Kansas City, Topeka, and St. Joseph pairs. Other metropolitan areas such as Pittsburgh and Cleveland are close enough to smaller neighboring stations like Youngstown to create the same overlap issue.

Overestimation, by contrast, can arise because of improvements in TV towers over time. In most cities, the 1950's saw expanded broadcast ranges through both upgrades to existing stations and also construction of new towers. The 2003 DMAs are therefore prone to overstate early 1950's signal reach, when towers were weaker. As shown in figure E.7, 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

¹¹Condition (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.

from 721 to 992 feet. Average visual power jumped from 19 to 170 kilowatts over that period, and average aural power increased from 11 to 87 kilowatts.¹² 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 E.5 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 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 propagation are discussed in section 2.1 of the main text. To validate the ITM measurements, we turn next to comparisons of key findings in the literature using the DMA approximation and ITM data.

D Empirical Appendix

D.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 personnel, and agricultural workers. To work with a consistent sample, we drop occupations that first receive coverage during this period.¹³ 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 30.¹⁴

¹²Power does not map directly to broadcast reach, as higher frequency channels require more power to operate.

¹³This 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.

¹⁴This drops an additional 1,996 individuals.

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 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.

D.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 E.1. 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 representativeness 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 [E.2](#) 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 [E.2](#) 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 [E.4](#) 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.

D.3 Robustness Checks

D.3.1 Stocks vs Flows

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

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

strong timing assumptions and it is unclear whether the AB estimator yields an unbiased estimate.

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

D.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.¹⁵ Specifically, we target population aggregates in an MSA, as well as their education and age demographics.

Table E.3 shows the baseline results with the weighted sample. The main takeaway is that the results are broadly similar to those reported in our baseline results (see Table 1). The point estimates are slightly larger for both groups and now also boarder line significant for the younger age group.

D.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 E.4, column 1). This suggests that the attenuation bias from mobility is relatively minor.

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

D.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

¹⁵We linearly interpolate values in between the 1950 and 1960 Census.


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 E.4 reports treatment effects for this non-complier group in column 1. Using $\sigma = 0.95$ and $NCTE = -0.224$ in the ATT formula yields an ATT of -0.273, very close to the ITT estimate of -0.270.

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


E Appendix Figures and Tables

Figure E.1: 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 E.2: 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 E.3: 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 E.4: 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 E.5: 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 E.6: 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 E.5. In Panel (B), for Minneapolis-St. Paul, “Coverage Map Counties” refers to those ringed in Figure E.5 (C), whose coverage exceeds 50 percent according to *TV Factbook* coverage maps. “Overreach Counties” refers to those highlighted in Figure E.5, which fall inside the Minneapolis-St. Paul DMA but outside the *TV Factbook* broadcast range.

Figure E.7: 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 E.1: Summary Statistics

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

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

Table E.2: Effect of Television on Employment Stocks and Flows

	(1)	(2)	(3)	(4)
		<i>Employment</i>		ΔEmp
Effect of TV stations $\times$				
Age over 50	-0.572*** (0.155)	-0.462*** (0.102)	-0.500*** (0.109)	-0.276*** (0.0683)
Age under 50	-0.270** (0.127)	-0.203** (0.0856)	-0.199** (0.0933)	-0.0911 (0.0590)
Female	0.187 (0.159)	0.127 (0.109)	0.157 (0.118)	0.0250 (0.0831)
Observations	292,448	292,448	260,895	292,448
R-squared	0.736	0.774	0.146	0.048
Year FE	Yes	Yes	Yes	Yes
Person FE	Yes	Yes	Yes	Yes
Specification	Stock	LDV	AB	Flow
Mean DV age over 50	52.01	52.01	52.24	0.340
Mean DV age under 50	56.11	56.11	56.36	1.420

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

Table E.3: Individual-level Effects of TV on Employment - Weighted Sample

	(1)	(2)	(3)	(4)	(5)
	$\Delta Employment$				
Effect of TV stations $\times$					
Age over 50	-0.308*** (0.113)	-0.323*** (0.114)	-0.323*** (0.120)	-0.339*** (0.126)	-0.320*** (0.121)
Age under 50	-0.277*** (0.104)	-0.275*** (0.104)	-0.241** (0.112)	-0.253** (0.117)	-0.240** (0.113)
Female	0.0526 (0.163)	0.0527 (0.163)	0.0458 (0.162)	0.0602 (0.190)	0.0459 (0.162)
Observations	288,249	288,249	288,249	288,249	288,249
R-squared	0.073	0.073	0.073	0.074	0.073
Year FE	Yes	Yes	Yes	Region X Year	Yes
Person FE	Yes	Yes	Yes	Yes	Yes
Age FE	No	Yes	Yes	Yes	Yes
Quartic Trends	No	No	Region	No	Demographics
Mean DV age over 50	0.350	0.350	0.350	0.350	0.350
Mean DV age under 50	1.422	1.422	1.422	1.422	1.422

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

Table E.4: Heterogeneous Effects of TV on Individuals Aged 50+ by Demographic Groups

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

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

References

- Aguiar, Mark, Mark Bilz, Kerwin Kofi Charles, and Erik Hurst**, “Leisure Luxuries and the Labor Supply of Young Men,” *Journal of Political Economy*, 2021, 129 (2).
- 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.
- 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.
- Gentzkow, Matthew**, “Television and voter turnout,” *The Quarterly Journal of Economics*, 2006, 121 (August), 931–971.
- **and Jesse M. Shapiro**, “Preschool Television Viewing and Adolescent Test Scores: Historical Evidence from the Coleman Study,” *Quarterly Journal of Economics*, 2008, 123 (February), 279–323.
- 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).
- Mcgrattan, Ellen R and Richard Rogerson**, “Changes in Hours Worked, 1950-2000,” *Federal Reserve Bank of Minneapolis Quarterly Review*, 2004, 28 (1), 14–33.
- Rachel, Łukasz**, “Leisure-Enhancing Technological Change,” *working paper*, 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).