

Growing Old:

How Ageing Slows Growth through New Product Adoption

Enrico D. Turri*

October 11, 2026

Abstract

This paper studies how ageing slows down growth by weakening incentives to innovate. I use scanner data to show that younger households are early adopters: relative to households aged 30, the share of spending going to new products is 29% lower at age 65 and 35% lower at 80. To study the implications of this fact for growth, I build a Schumpeterian model of innovation in which agents become familiar with products they consume. Older consumers spend less on new products, being more familiar with existing ones. I show that ageing decreases both creative destruction and growth, since new products face reduced sales to older buyers. Calibrating the model to the US, I find that ageing explains 32.1% of the GDP growth slowdown between the 1980s and the 2010s. Along the transition path, ageing generates a growth expansion before the eventual slowdown. While ageing weakens innovation, it increases welfare and makes research subsidies more effective. As innovators ignore the familiarity new products destroy, the model uncovers a new externality that biases growth upwards in the decentralised equilibrium.

Keywords: Population ageing; new product adoption; early adopters; consumer familiarity; firm entry; creative destruction; economic growth; innovation externalities; research subsidies.

JEL Codes: D12, J11, O31, O33, O41.

*London School of Economics. Email: e.turri@lse.ac.uk. I'd like to thank Ethan Ilzetzki, Silvana Tenreyro, and Pietro Garibaldi for their continuous encouragement and support throughout my PhD studies. I am indebted to Maarten De Ridder for detailed comments and suggestions. I also thank Daron Acemoglu, Matthias Doepke, Thomas Drechsel, Steve Pischke, Jane Olmstead-Rumsey, Iván Werning and seminar participants at LSE and MIT for comments and suggestions. Researcher(s)' own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

1 Introduction

The economies we live in are ageing. In the United States, the median age of the population rose by almost ten years between 1980 and 2025, from 30.0 to 39.4 years of age. Over the same period, the share of Americans aged 65 and over grew from 11.3% to 18.9%, and the Census Bureau projects that it will reach 24.4% by 2060 (Figure 1a). Around the world, the US is not an exception: Japan and much of southern Europe have aged even faster, and so have large emerging economies. In Brazil and China the median age almost doubled in the last forty years (Figure 1b). This demographic shift is unprecedented in human history (United Nations, 2002), and raises a question: what does population ageing imply for innovation and economic growth?


This paper argues that ageing slows down growth because an older society spends less on new products, and that this in turn weakens incentives to innovate. Innovation drives growth: new products are of higher quality than the ones they replace, or cheaper to produce (Aghion and Howitt, 1992), and are brought to the market by potential innovators who weigh the cost of research and development against the profits from future sales. Older consumers, however, are slower to adopt new products and spend less on them, preferring the ones they already know. As the population ages, new products sell less and expected profits fall. Faced with lower profits, innovators invest less in research, and economic growth slows down.

I analyse the effect of population ageing on growth in three steps. First, I use scanner data to document that older consumers spend less on new products. Older consumers are commonly thought to be late adopters, but the evidence in the marketing and sociology literatures is mixed.¹ I use the NielsenIQ Homescan Consumer Panel, which records the purchases of about 50,000 US households a year between 2004 and 2024 and I identify 111,002 new product launches which I use in my empirical analysis. I find that older households devote a smaller share of their spending to new products: relative to households aged 30, the share is 15% lower at age 50, 29% lower at 65 and 35% lower at 80. The spending gap is robust to controls for income, education, household composition, as well as for shopping behaviour and for the categories of products older households buy.

Motivated by this evidence, in the second part of my analysis I develop a Schumpeterian

¹The study of the diffusion of innovations was pioneered by rural sociologists (Ryan and Gross, 1943), and the “early adopter” category was introduced by Rogers (1958). Rogers (1983, p. 251) concludes from a collection of 228 studies that earlier adopters are not different from later adopters in age. Using household purchase data on many launches, Steenkamp and Gielens (2003) find no general effect of age on the trial of 239 new products in the Netherlands, while Gielens and Steenkamp (2007) find lower first-year purchases among older consumers for 301 new products in four European countries. The marketing meta-analysis by Arts et al. (2011) finds that demographics explain little of the variation in adoption. Lambert-Pandraud et al. (2005) and Lambert-Pandraud and Laurent (2010) show that older buyers repeat purchases of established brands of cars and perfume. In economics, Bornstein (2026) uses scanner data from the same dataset I use to show that older households switch less often between existing brands.

Figure 1: Population ageing, 1980–2025



(a) Median age of the US resident population (left axis) and share of the population aged 65 and over (right axis), July 1 of each year, from Census Bureau national population estimates. (b) Median age in 1980 and 2023 in the G7, in the most populous country of each continent and in the world, from the estimates of the United Nations' World Population Prospects 2024, whose last year is 2023. The UN estimates put the US median age about one year below the Census figures of panel (a).

model of innovation that matches the life-cycle pattern of spending on new products documented in the data. The mechanism that generates it is that consumers become familiar with individual products over time.² Older consumers have had more time to become familiar with existing products, so they like them better and spend less on new ones. Younger consumers, less attached to continuing products, spend a larger share of their budget on new ones instead. The model features heterogeneity on both the household side, where agents differ in age, and the product side, where products differ in quality. Over time, product quality increases through innovation, which is carried out by entrants investing in research and development motivated by the prospect of monopoly rents (Aghion and Howitt, 1992). When they are successful in innovating, entrants replace incumbents in the process of creative destruction. A new product is thus of higher quality, but every consumer starts out unfamiliar with it.

In the model, ageing lowers innovation and growth. All else equal, ageing does not reduce total spending: it shifts sales from young products to old ones. Older consumers buy less of a

²Familiarity has a long tradition in psychology. In the mere-exposure effect tradition, repeated exposure to a stimulus is enough to make people like it more (Zajonc, 1968), and a meta-analysis of 208 studies confirms the effect (Bornstein, 1989). In economics, Bronnenberg et al. (2012) use scanner data to show that brand preferences formed through past consumption persist for decades.

new product at launch, but as time passes and they become familiar with it, they spend more on the product. Innovators, however, care about the net present value of sales, so the losses at launch outweigh the later gains. The value of entering then falls, and creative destruction falls and growth slows down. Additionally, I show that moving expenditure from younger to older consumers has the same effect as population ageing, as the expenditure-weighted consumer becomes older. Thus both population ageing and the reallocation of expenditure to older households contribute to the growth slowdown.³

In the third part of my analysis, I calibrate the model to the United States to quantify the effect of ageing on growth, and show that ageing explains a sizeable share of the growth slowdown that occurred over the last four decades. In the calibration I use moments from both macroeconomic aggregates and microdata: I discipline how households' familiarity with products evolves over time so that the model matches the age profile of spending on new products documented in the first part of the analysis. Comparing decades, the ageing observed between the 1980s and the 2010s lowers growth in the model from 2.10% to 1.95% a year, while in the data growth fell from 2.10% to 1.64%: ageing explains 32.1% of the growth slowdown.⁴

The model delivers further insights into the mechanisms through which ageing affects growth. Along the transition, ageing raises growth before lowering it. Since the age distribution and the familiarity of consumers change slowly, an innovator who enters before the population has aged sells at launch as much as before, but expects to sell more later, once its product is old and consumers have aged. The value of entering early therefore rises. In the simulation from 1980, growth rises from 2.02% to 2.05% a year in the early 1980s, and stays above its initial level until 1990. In these years growth looks healthier than before, while the fundamentals of innovation have already deteriorated. Two additional quantitative results follow from the transition exercise. First, had the US kept its 1980 demographics, its GDP per capita would have been 3.5% higher in 2025. Second, even if ageing were to stop in 2025,⁵ growth would not recover: it would stay at about 1.85% a year over the coming decades, well below the 2.02% of 1980. The gap with the economy that had not aged would then reach 8.0% by 2050.

Ageing also makes research subsidies more effective. A subsidy that covers 10% of the cost of research raises growth by 0.42 percentage points in the economy of the 1980s and by 0.47 in that of the 2010s, 11.0% more. In an older society, entrants face a larger stock of familiarity with

³In the US, between 1980 and 2025, the mean adult age weighted by expenditure rose by 6.8 years, against 5.4 years for the population. In particular, by the end of the 1980s the spending of older people had caught up with that of younger people (Gokhale et al., 1996).

⁴Population ageing alone, assuming households didn't differ in spending, explains 11.1% of the growth slowdown, with the remaining share coming from the reallocation of expenditure towards older households that happened in the same period.

⁵The Census Bureau projects that ageing will continue at least until the end of the century. I use Census projections in an alternative simulation of my model.

existing products. Higher creative destruction reduces this stock, since faster product churn leaves households less time to become familiar with each product. Each additional entrant thus makes it easier for the following ones to sell at launch, and in an older society this effect becomes stronger, while the higher risk of being replaced due to higher creative destruction changes by less between the two societies. The value of entering therefore falls more slowly as entry rises in an older society, and the same subsidy buys more innovation.

Despite slowing growth, ageing raises welfare. Two channels specific to the model counter-balance the loss from slower growth: longer lives give consumers more time to become familiar with products, and lower creative destruction destroys less of this familiarity. In the calibrated model, these gains outweigh the growth loss.⁶ Comparing balanced growth paths, ageing raises welfare by 3.3% in consumption equivalents.⁷ The lifetime utility of a newborn rises more, by 6.2% along the observed demographic path, mostly because she values living longer.⁸

Finally, the model uncovers a new externality. Innovators weigh their own expected profits against the cost of research, but ignore that a new product destroys the familiarity consumers had built with the one it replaces. A social planner takes this loss into account instead. I call this externality *familiarity destruction* and show that it pushes the decentralised economy towards too much creative destruction. This externality can change the verdict of the classic Schumpeterian wedges. In the calibrated model, considering only appropriability, intertemporal spillovers and business stealing, the planner would want more innovation in both decades at the market's rate of entry; taking familiarity destruction into account, the planner values one more innovation at only 68% of its cost in the 1980s and 53% in the 2010s, and would thus reduce innovation, even more in the older society.⁹

In the final part of the paper, I show evidence in support of four implications of the model. First, R&D spending should fall in consumer-facing industries as the population ages: between 2013 and 2023, the R&D of the industries that make packaged consumer goods fell by 9% in

⁶This result rests on familiarity raising utility. If familiarity only shifted demand without enhancing utility, for example because consumers know less about new products, the familiarity gain would disappear and households would misallocate their spending towards the products they know. Ageing would then lower welfare, by 0.9% along the observed demographic path and by 2.6% across balanced growth paths (Section 7.3).

⁷The model is well suited to this exercise because preferences do not change with age: older consumers differ from younger ones only in the familiarity they have accumulated. Along the observed demographic path, the gain is 1.6%, because the initial increase in growth destroys familiarity.

⁸Social welfare does not value longer lives as such, since each household that dies is replaced by a newborn and the total population is constant.

⁹As for welfare, this result hinges on familiarity raising utility. If familiarity only shifted demand, destroying it would reduce the misallocation of spending towards the products households know, and the planner would value one more innovation at 166% of its cost in the 1980s and 175% in the 2010s: the market would innovate too little (Section 8). The size of the externality also depends on the assumption that a new product drives the old one out of the market. In reality, old products need not exit, and the externality is smaller when they survive. Many products do exit, however, and consumers miss them.

2023 dollars, while that of all US businesses rose by 73%. Second, markets whose buyers are older should see fewer new products. Across the product modules of the HMS, a module whose buyers are ten years older launches 14% fewer new products per product on sale. Third, young firms should be smaller relative to mature firms where buyers are older. I use Compustat and the CEX to show that among listed firms in 47 consumer industries, the sales gap between firms in their first five years and firms more than ten years old widens by 0.64 log points per ten years of buyer age. Fourth, in my model the price-earnings ratio of a firm falls with its age, a result that quality-ladder models fail to produce, and ageing flattens its decline. In Compustat the ratio falls by 0.44 log points over the forty years after listing, and the fall has shrunk by more than half between the 1980s and the 2010s, as the buyers of US firms aged.

Related literature. This paper contributes to three strands of literature. The first is the one on the macroeconomic effects of ageing. Here, Bornstein (2026) uses the HMS to document that older consumers switch less often between existing products, and studies how this inertia negatively affects firm entry using a model with stochastic switching costs. My work instead builds a new, tractable model of Schumpeterian growth that features ageing and demand frictions, which allows me to quantify the consequences of ageing for innovation and growth and to conduct a normative analysis. In related work, Dang (2025) shows that the interaction of customer habits with ageing widens the innovation gap between market leaders and followers, generating rising concentration and markup dispersion. Brès and Angelini (2025) document the effect of ageing on competition and use a model of search in the product market to study the effect of ageing on the level of GDP. Crouzet et al. (2026) show that in India younger consumers are more likely to use mobile payments, and studies how demographics interacts with network externalities. Acemoglu and Linn (2004) show that demographic change shapes innovation through market size: as the age structure of the US population changes, more new drugs enter the categories whose potential market grows. In my paper, ageing instead shifts spending from new to old products within the same markets. There is also a large literature on how ageing affects firm entry and growth by reducing the supply of young workers. Karahan et al. (2024) show that the slowdown in the growth of labour supply explains a sizeable share of the secular decline in the US start-up rate, and Hopenhayn et al. (2022) that falling population growth accounts for the shift of employment towards older firms, rising concentration and the falling labour share. Peters and Walsh (2026) study the consequences of declining population growth in a model of firm dynamics and show that it weakens economic growth. Engbom (2019) obtains a similar result by combining idea growth and matching in the labour market. Lehr (2026) studies the age of workers instead: younger workers adopt new technologies more, and a younger workforce raises technology adoption, wages and innovation. In his model, the baby boom accounts for

about a quarter of the acceleration of US productivity growth in the late 1990s. While in these papers demographics enter the picture through the supply channel of the size, growth and age of the workforce, this paper studies a complementary consumer demand channel. Finally, work on automation finds that ageing can increase growth by inducing the adoption of labour-saving technologies (Acemoglu and Restrepo, 2017, 2022; Acemoglu et al., 2026).

The second strand of literature to which my paper contributes is the one on customer capital (Gourio and Rudanko, 2014), of which familiarity with products is one specific form operating through good-specific habits (Ravn et al., 2006). Bronnenberg et al. (2012) use the HMS to show that this capital is built by past consumption and persists for decades, creating large entry barriers for new brands. In existing models within this literature, most recently Ignaszak and Sedláček (2026), the customer base is an asset of the firm, built by spending on customer acquisition or by reducing markups right after entry to encourage early sales, only to raise them once the customer base is built. Unlike in these papers, firms in my model cannot influence familiarity growth by strategic investments or pricing. I contribute to this literature by introducing the demographic dimension into it.

Finally, the paper contributes to the debate on the causes of the productivity slowdown (Fernald, 2014) and of declining business dynamism (Decker et al., 2014) by identifying a new channel that potentially explains part of the observed slowdown in growth. Mechanisms proposed in the literature include a decline in knowledge diffusion from frontier to laggard firms (Akçigit and Ates, 2021, 2023), intangible inputs and falling costs of scale that entrench market leaders (Aghion et al., 2023; De Ridder, 2024), more incremental innovation by smaller firms (Olmstead-Rumsey, 2026), ideas becoming harder to find (Bloom et al., 2020), low interest rates that strengthen market power (Liu et al., 2022), and the diversion of innovation towards leisure-enhancing technologies (Rachel, 2024). Ageing provides a demand-driven explanation for part of the growth slowdown.¹⁰

Outline. The paper proceeds as follows. Section 2 studies the demographics of new product adoption in the data. Section 3 introduces the model. Section 4 calibrates the model to the US and quantifies the impact of ageing on growth comparing the 1980s and the 2010s. Section 5 simulates the transition using the observed demographic path. Section 6 explores how ageing changes the effectiveness of research subsidies. Section 7 studies welfare and shows that ageing it increases along with individual utility, even though it slows growth. Section 8 characterises

¹⁰Cravino et al. (2022) perform a VAR analysis to study the effect of ageing on macroeconomic aggregates, finding a negative effect on growth. They explain this result with a model in which innovation depends negatively on the old-age dependency ratio. Maestas et al. (2023) exploit differences in the pace of ageing across US states and find that ageing lowered the growth of GDP per capita by 0.3 percentage points a year between 1980 and 2010, two thirds of it through slower growth of labour productivity.

the first best and the familiarity-destruction externality. Section 9 takes three untargeted predictions of the model to the data. Section 10 concludes.

2 The demographics of new product adoption

This section studies how spending on new products changes with age, using scanner data from the NielsenIQ Homescan Consumer Panel (HMS). It first describes the most salient features of the HMS, how I define a product and how I identify new product launches, and compares new products with the other products in the panel. It then performs the main empirical exercise of the paper estimating the age profile of the share of spending that goes to new products, which is later used to calibrate the model.

2.1 Dataset and new products

The main features of the dataset are summarised in Table 1.

Table 1: The Homescan panel

	NielsenIQ Homescan
Structure	panel; median tenure: 3 years; mean tenure: 5.4 years
Years covered	2004–2024
Unit of observation	household $\times$ year
Households per year	37,786–63,350
Household-years, total	1,203,446
Coverage	barcoded packaged goods; random-weight items excluded
Expenditure recorded	\$4,609 per household (2024)
Other covariates	income, education, size, children, ZIP code

The HMS records purchases from 2004 to 2024 for a rotating panel of between 38,000 and 63,000 households a year. The median and mean tenure of panellists are 3 and 5.4 years, respectively. The households in the panel are instructed to record every purchase of products that have a barcode, using a barcode scanner provided to them.¹¹ These products are typically

¹¹Barcodes, officially known as Universal Product Codes (UPCs), were introduced in the early 1970s, with the first commercial scan happening on 26 June 1974 in Ohio. They identify each product with a sequence of twelve digits, of which the first six to nine form the company prefix and the next two to five digits identify individual products. The twelfth and final digit (the check digit) is mathematically calculated based on the previous numbers to ensure the barcode scans accurately. Each digit is then represented by a fixed-width pattern of alternating light spaces and dark bars. UPC company prefixes are assigned by GS1, a global non-profit standards organisation. By the mid-to-late 1980s, UPC barcodes had become a market standard. Economic research has shown that scanners increased a store's labour productivity by 4.5% in the first few years after their introduction (Basker, 2012) and stimulated variety-enhancing product innovation and the growth of international retail supply chains (Basker and Simcoe, 2021).

found in supermarkets and similar retail outlets. All barcodes that appear in the HMS are classified, in increasing detail, into 10 “departments” (e.g., dry grocery, frozen foods, health and beauty care), 118 “product groups” (e.g., carbonated beverages, pet food, snacks) and 1,170 “product modules” (e.g., carbonated soft drinks, refrigerated yogurt, dry dog food).¹²

The HMS provides age information on all members of the household and identifies one or two heads, at most one male and one female.¹³ My main age measure is the age of the younger head. I use head age since heads are more likely to be the ones who choose which products the household buys. Alternative choices for age when more than one head is present, such as the age of the older, male or female head, or the average head age, do not affect the sign or order of magnitude of the results reported in Section 2.2 or of the calibration. In my analysis, I also show that results broadly survive when restricting the sample to households made up of a single person (in fact, I find most age effects to be stronger in this setting). The average age of all individuals in all households in the survey is only a couple of years above the average age in the population as reported in the Census. Head age, instead, is about thirteen years higher. Both age measures increase over time along a trend parallel to that appearing in the Census.¹⁴ Additional information on household characteristics includes income, education and ZIP code of residence. NielsenIQ also supplies household projection weights to make the sample representative of the US population. Given its content and its panel structure, this dataset is particularly suited to the study of how age affects purchasing choices, since it allows one to map individual product purchases to households and to follow both households and products over time.

The panel records spending of several thousand dollars per household each year, \$4,609 per household in 2024. A natural question is what fraction of total household spending is captured by the HMS. Using the CEX to measure the universe of household spending, Jaravel (2019) finds that between 2004 and 2015 spending recorded in the HMS was equal to 6.7% of total household spending. This share derives both from the HMS recording purchases for selected categories only, and from panellists not scanning all their purchases within these categories. In the CEX, food, drinks and household supplies, which the HMS covers, account for 8.6% of total spending.¹⁵ In addition, goods only account for around 34% of spending, so the share of goods expenditure recorded by the HMS is around one fifth. Jaravel (2019) also finds that the HMS

¹²These are departments, groups and modules with at least one purchase by a panellist over the dataset period 2004–2024; NielsenIQ’s documentation lists a few hundred additional modules that never appear in the panel’s purchases.

¹³Heads are self-reported by the household and two heads need not be a couple. Panellists report at sign-up whether the household has one head or two, and NielsenIQ’s documentation says that the two heads “could be married, single, related, etc”. A mother and her adult son can be the female and male heads. The documentation is silent on how heads are identified in same-sex couples.

¹⁴See Figure 15a in the Appendix.

¹⁵HMS goods account for 13.4% of spending once goods that it covers only in part, such as apparel and personal care products, are added.

expenditure share falls with income, consistent with Engel’s law, since spending on food and other necessities, which make up most of the HMS, grows less than total spending.

In the context of this dataset, I define a product as a brand-module pair. The brand is the name consumers recognise and become familiar with, while the module separates different goods sold under the same name. For example, “Cheerios” ready-to-eat cereal is one product, although it is sold under 153 barcodes in 33 package sizes. “Dove” bar soap and “Dove” ice cream bars, instead, share a brand name but are two different products. The same unit is used by Michelacci et al. (2022), who define a variety as a brand-module pair. They find that, when varieties are identified by barcodes, most of the contribution of new varieties to fluctuations in household purchases comes from new barcodes of existing brands, which they interpret as a form of strategic marketing. Studies based on NielsenIQ data usually define a product as a barcode (Broda and Weinstein, 2010; Jaravel, 2019; Neiman and Vavra, 2023; Argente et al., 2024; Einiö et al., 2025; Argente et al., 2025), sometimes using brands as a robustness check. A new barcode, however, is often a new package or a minor variant of a product that consumers already know. Broda and Weinstein (2010) document that most barcode entry and exit takes place within existing brands, and Argente et al. (2024) find that sales of brands decline more slowly over their life cycle than sales of barcodes, as firms renew existing brands with new barcodes. Bornstein (2026) instead uses a coarser unit and defines a brand as all barcodes sold by a firm within a module. Under this definition, a new brand launched by a firm that already sells in that module is not a new product. Among all brand-module debuts, I then select a subset of debuts that introduce a genuinely new product, as detailed below. Across the 21 years of the panel, according to my definition, the median and mean number of products available in each year are 125,484 and 120,532, respectively. On average, in each year a household purchases 370 distinct products (with a median of 351).

I identify the date on which a brand-module pair is launched on the market as the one in which it is first purchased by at least one panellist. In total, the dataset features more than four hundred thousand debuts. However, many of these do not represent true new product introductions. First of all, for all products present in the first year of the panel, I cannot tell whether they have just been introduced or predate the beginning of the data collection, so I exclude them. The same holds for products sold in the first year in which a module is introduced.¹⁶ I then exclude all private labels (retailer store brands, as flagged by NielsenIQ) and pseudo-brands.¹⁷ I also exclude manufacturer names that come without a brand, which

¹⁶A total of 155 modules are introduced between 2007 and 2024, including, for example, “Fish care” and “Pet collars and leashes”. Most products sold in newly introduced modules were already in the dataset the year before, with respectively 93% and 71% of first-year sales coming from previously existing barcodes.

¹⁷Pseudo-brands are “GENERIC”, “NO COMPANY LISTED”, “MANUFACTURER NOT AVAILABLE”, “NBL-NO COMPANY LISTED” and “MULTIPLE VALUE”.

NielsenIQ marks as “NBL” (no brand listed). In addition, some brands change their names (for example, PepsiCo renamed “Mountain Dew” to “Mtn Dew” in January 2009) or are reallocated by NielsenIQ to a different module for no apparent reason related to the product itself (brand names are assigned by NielsenIQ; see the discussion in Neiman and Vavra, 2023). To address this issue, I exclude products that draw part of their sales from barcodes that were active in the previous two years under a different brand name or module.¹⁸ This leaves me with a total of 280,035 product launches. Finally, I exclude brand extensions, in order to capture original new products coming to the market.¹⁹

I identify 111,002 new product launches in the HMS dataset between 2005 and 2024, for an average of 5,550 launches per year. Details on how many brand-module launches are excluded at each of the steps described above are in Table 2.

The number of launches per year shows no clear trend: it moves between 4,134 (2024) and 6,953 (2013).²⁰ The peak year in terms of sales is the launch year for 40% of launches and the following year for another 20% of them.²¹ Sales are concentrated: the top decile of launches accounts for 81% of expenditure on newly introduced products. Launches happen in almost all modules in the dataset (1,049 out of 1,170), but are distributed unevenly across departments. Dry grocery accounts for 28% of launches and health and beauty care for 15%. Alcoholic beverages, mostly wine and beer, account for a quarter of launches but only 11% of spending on new products in the launch year. For details, see Table 14 in the Appendix.

Table 3 compares new products with brand extensions launched in the same years and with continuing products already on sale in 2004. New products are launched throughout the year, with the median launch in mid-June. They sell to few households at first: the median new product reaches 2 panel households in its launch year, against 12 for a continuing product in a typical year. They also persist less on the market: 62% of new products are still bought one year after launch and 22% ten years after, against 86% and 52% of continuing products over

¹⁸Because a product is a brand-module pair, it can sell under more than one barcode at the same time. An example is a detergent sold in packages of different sizes. In the Mountain Dew case above, more than 97% of first-year sales came from previously active barcodes and the rename is excluded from the pool of launches.

¹⁹Brand extensions take two forms. In the first, the brand name has already recorded sales in a different module. In the second, the brand name is a “name child” of another brand: there exists an older brand name, either in the same or in another module, the new brand’s name is a continuation of, as “Chobani Flip” is of “Chobani”. A single-word name, however, can be a common word shared by unrelated brands. When the older name is a dictionary word or a place name, I count the new brand as its name child only if at least a quarter of the older brand’s name children sell in modules where the older brand had already sold, counting under the older brand the names to which NielsenIQ adds a suffix, such as “Tide - H-D Liq”. For example, “Hello Bello” diapers are a new product and not an extension of “Hello” toothpaste. A leading run of initials, such as “G M” for General Mills, counts as a single word.

²⁰See panel (a) in Figure 16 in the Appendix.

²¹This early peak result is consistent with Argente et al. (2024), who document, using mainly NielsenIQ retail scanner data, that sales rise only briefly after a new product is introduced, and then decline steadily.

Table 2: What counts as a new product in the Homescan panel

	Brand-modules	Removed
All brand-module first appearances	429,257	
less already selling at panel start	323,264	105,993
less module's own first year	321,161	2,103
less private label	320,944	217
less generic pseudo-brands	320,027	917
less manufacturer names without a brand	308,496	11,531
less debuts on recently relabelled barcodes	280,035	28,461
less brands already selling in another module	168,918	111,117
less name extensions of an older brand	111,002	57,916
Launch sample	111,002	

Notes. Brand-module pairs at their first purchase by a panellist, 2004–2024. Each row applies one of the exclusions described in the text to those left by the row above, and the *Removed* column counts the number of brand-modules the exclusion drops.

Table 3: New products, brand extensions and continuing products

	New products	Brand extensions	Continuing products
Number of products	111,002	169,033	104,687
Panel buyers in the launch year (median)	2	2	12
Spending in the launch year, \$ thousand (median)	30.0	33.2	171.3
Launch day of the year (median)	June 17	June 4	–
Still bought 1 year later (%)	61.8	66.6	86.3
Still bought 5 years later (%)	33.7	38.7	67.7
Still bought 10 years later (%)	22.0	25.1	51.6
First-time buyers who buy again the next year (%)	16.3	18.2	19.6
Unit price relative to the module median (%)	+20.8	+14.7	–6.8

Notes. New products are the 111,002 launches of 2005–2024. Brand extensions are the brand-modules launched in the same years that are excluded as extensions in Table 2. Continuing products are the branded brand-modules already on sale in 2004; for them, launch-year rows refer to a typical year between 2005 and 2024, and survival to products sold in a given year. Spending uses projection weights. Repeat purchase: households buying the product for the first time in a year, present in the panel the year before and the year after, who buy it again the following year. Unit price: price per unit of volume, relative to the median branded product in the same module, size unit and year, in the launch year for new products and brand extensions.

the same horizons. Their buyers also buy them again slightly less often: among households buying a product for the first time, 16% buy it again the following year when the product is new, against 20% when it is a continuing one. Finally, new products are more expensive: their median unit price is 21% above the median branded product in the same module, while the median continuing product is 7% below it. Brand extensions mostly sit between the two groups.

On the household side, the average household adopts 57.8 new products in its spell as a panellist. The 3,496 households who never buy any of the new products make up 1.6% of all households who ever appear in the dataset, while the household with the most new product adoptions purchases 1,698 new products over fourteen years in the panel.

2.2 Age and expenditure on new products

Regression specification. To measure how spending on new products changes with age, I estimate, for household i in year t ,

$$s_{it} = \alpha_t + \beta_{\text{age}(i)} + \gamma' X_{it} + \delta_{c(i)} + \zeta' Z_{it} + \eta E_{it} + \varepsilon_{it}, \quad (1)$$

where the dependent variable s_{it} is the share of spending that goes to products launched in year t . Using the share rather than the amount spent removes differences in total spending between younger and older households. The coefficients of interest are the age effects $\beta_{\text{age}(i)}$, where $\text{age}(i)$ is the age of the younger head in year t . There is one age effect for each age between 26 and 85, and one for each of the two groups aged 25 and below and 86 and above, with $\beta = 0$ at age 30, so that each effect is the difference from households aged 30. The calendar-year fixed effects α_t absorb effects common to all households in a year, such as the number of products launched each year and their success. I estimate the equation by least squares with NielsenIQ's projection weights, and cluster standard errors by household. The sample includes every year in which a household was in the panel between 2005 and 2024: 1,163,731 household-years from 219,109 households.

I include several controls to isolate the association between age and spending on new products. First, older households differ from younger ones in income, education and household composition, and each of these can change how much a household spends on new products. The household controls X_{it} include indicators for 16 income brackets, for six education levels of the younger head and for education not stated, for the presence of children under 18, and for household size, with one, two, three, four, and five or more members. Older households may have also been in the panel for longer. Since households may record their purchases differently when they have just joined, X_{it} also includes indicators for the household's first and second

years in the panel.²²

Any difference between younger and older households could also come from differences between generations rather than age. Households born more recently may have grown up with more new products and have developed a stronger love for novelty. In this case, today's young would keep buying new products when old, and an older population would not spend less on them. The birth-year effects $\delta_{c(i)}$ allow us to compare households born in the same year at different ages. Since age equals calendar year minus birth year, age, year and birth-year effects cannot be separately identified. I use the Deaton–Paxson restriction (Deaton, 1997), applied to the birth-year effects: they sum to zero and have no linear trend.²³

Older households may also shop more carefully. They have more time, especially after retirement, and can use it to look for deals and cheaper stores (Aguiar and Hurst, 2007). A careful shopper may avoid new products, which cost more than the products already on sale (Table 3). To rule out this alternative explanation, the terms Z_{it} measure employment and shopping behaviour in year t and include

- for the male and the female head, indicators for working part time, not working, and no such head, with full time as the omitted category;
- for the male and the female head, an indicator for being 65 or older and its interactions with the employment indicators;
- the logs of the number of shopping trips, of stores and of retail chains visited;
- the share of spending bought on deal;
- the share of spending saved with coupons;
- the share of spending on store brands;
- a price index, the ratio between what the household paid and what it would have paid at the median price of the same barcodes in the same month and state, as in Aguilar and Hurst (2007);
- the shares of spending at discount stores, warehouse clubs, dollar stores, drug stores, on-line and other retailers, with grocery stores as the omitted channel;


²²I count the years in which the household was in the panel since 2004, the first year of the data. The panel started before 2004, so households already in it in 2004 may have joined earlier, and they indeed seem to behave like long-standing panellists: 79% of them are still in the panel in 2005, while only 65% of the households that join later stay beyond their first year. For this reason, they have a separate indicator for their second recorded year.

²³Birth years before 1922 and after 1997, which have few households, are pooled into one group at each end.

- the logs of the number of distinct products and of product modules bought, and of total spending.

Finally, older households may buy more often products belonging to categories in which fewer new products are launched. Similarly, they might have a preference for cheaper products, while new products are more expensive than continuing ones. They would then spend less on new products only because fewer are offered within what they usually buy. To address this confounding factor, I add the exposure to new goods $E_{it} = \sum_c s_{ict} N_{ct}^{(-i)}$ as a control that measures how much a household would spend on new products if, in each category and price tier, it behaved like all other households in the panel. The cells c split each product module into three: products below and above its median unit price, and products without a measurable unit price. Here s_{ict} is the share of the household's spending on products already on sale that falls in cell c in year t , and $N_{ct}^{(-i)}$ is the share of the cell's spending in year t that goes to products launched in t , among all other households.

Figure 2: Share of spending on new products by age



Age effects β of equation (1), estimated with all its terms: calendar-year fixed effects, household controls X_{it} , employment and shopping behaviour Z_{it} , exposure to new goods E_{it} and birth-year effects $\delta_{c(i)}$, with the Deaton–Paxson restriction (Deaton, 1997). One indicator for each head age between 26 and 85, and two pooled indicators for ages 25 and below and 86 and above (hollow markers), relative to age 30, the omitted category. The right axis reports the effects as a percentage of the average spending share of households aged 30, the left axis the implied spending share, that average plus the age effect. 95% confidence bands; projection weights; standard errors clustered by household; 1,163,731 household-years.

Results. Figure 2 shows that the share of spending on new products falls with age by plotting the age effects from regression (1). Households aged 30 devote 0.352% of their spending to products launched in the same year. Until the early forties, the share is not statistically different from this level. Then it falls steadily: relative to age 30, it is 14.6% lower at 50, 29.0% lower at 65

and 35.3% lower at 80, and it is lower at every age from 43 on. Between 40 and 80, each year of age lowers the share by about 0.81% of its level at 30, or 8.1% every ten years, and the decline is roughly linear in age. For innovators, an older population thus means lower sales in the year when the product is launched and a back-of-the-envelope calculation shows that the effect is sizeable. Between 2005 and 2024, US households spent on average \$1.83 billion a year on the new products I identify (in the products' launch years). If every household were ten years older, these sales would be 9.0% lower: \$165 million less a year.²⁴ The age profile obtained in this regression is the main empirical moment of the paper: in Section 4.2, I use it to discipline how fast consumers become familiar with products in the model.

Other controls. Table 15 in the Appendix reports the coefficients of the other terms included in the regression. Richer and more educated households spend a larger share on new products. Relative to an income of \$50,000 to \$60,000, the share is 0.044 percentage points higher above \$100,000, 13% of the level at age 30, and 0.035 points lower below \$5,000. College and post-college graduates spend 0.014 and 0.013 points more than high-school graduates. Children lower the share by 0.027 points, and households of four members or more by 0.013 to 0.017 points. A potential reason could be that parents stick to products the children already know, since a new product that the children refuse goes to waste. Careful shoppers, if anything, buy more new products, not fewer. Households that make more trips to the store and save more with coupons spend more on new products: twice as many trips raise the share by 0.010 points, and ten percentage points more of spending saved with coupons by 0.030 points. Households that pay more than others for the same barcodes spend instead less: paying 10% more than the median price lowers the share by 0.020 points. Deals and employment status have no effect. Exposure to new goods has a strong effect: a household whose categories and price tiers devote 0.1 points more of their spending to new products spends 0.039 points more on them, 11% of the level at age 30.

One-person households. The empirical exercise that was the subject of this section measured the age of a household with the age of its younger head. However, in a household with several members, the products bought might reflect the tastes of members of different ages, so the age of the head is only a proxy for the age of the buyer. Even though the main specification controlled for the presence of children and the household size, it is natural to ask whether

²⁴I use all household-years in the panel between 2005 and 2024, weighted with NielsenIQ's projection weights, which sum to the number of US households, and express spending in 2024 dollars with the consumer price index. In the counterfactual, each household keeps its total spending, and its share of spending on new products changes by the difference between the age effect of Figure 2 at its age plus ten years and the one at its age. The HMS covers only packaged goods, so these figures understate spending on new products in the whole economy.

the age gradient would survive if we were able to observe the purchasing choices of individual consumers. Figure 3 provides a definitive answer by repeating the estimation on one-person households. While restricting to this set of panellists lowers the sample size, this is still large and has 298,825 household-years from 58,737 households. In this specification, the controls exclude children and household size, but add a gender indicator. The decline is at least as steep as for all households: relative to age 30, the share is 18.7% lower at 50, 32.3% lower at 65 and 35.5% lower at 80.²⁵ The age gradient is therefore not a product of household composition: it is at least as strong when the age is that of the only person choosing which product to buy.

Late adoption. The analysis so far showed that older households spend less on new products. However, when an innovator considers introducing a new product on the market, she cares about the present value of all profits she will earn, not only sales in the first year. Figure 4 widens the horizon by replacing the spending share on new products with the share of spending on products launched in the last three years and in the last six years. All controls remain the same with the exception of the exposure E_{it} , which here is modified accordingly: it measures the exposure to products aged up to two and five years, respectively. Since the first launches in the data are from 2005, these two samples start in 2007 and 2010. The age profile is flatter for older products. Relative to age 30, the share of products aged 0 to 2 is 13.2% lower at 65 and 21.5% lower at 80, and that of products aged 0 to 5 is 10.1% and 16.0% lower, against 29.0% and 35.3% for new products. Older households are late adopters and catch up in part as products age, but the spending-share gap with younger households does not close within five years.

3 A Schumpeterian model with ageing and adoption


This section develops an overlapping-generations model of Schumpeterian growth with perpetual youth, in which consumers enjoy higher utility from products they are familiar with. The section ends with the characterisation of the balanced growth equilibrium under a fixed parameter configuration. How ageing affects the economy is then the subject of the following section.

3.1 Demographics, products, and familiarity

Time is continuous, indexed by $t \in \mathbb{R}$, and the economy is populated by a positive measure of households indexed by i . Each household faces a constant Poisson hazard of death $\mu > 0$,


²⁵ Estimates are less precise compared to the full sample, especially since few young people live alone: the sample has only 1,180 household-years at age 30.

Figure 3: Share of spending on new products by age, one-person households



Age effects β of equation (1) estimated on one-person households, with all its terms as in Figure 2, except that the household controls X_{it} exclude children and household size, and the employment terms in Z_{it} are an indicator for a woman, indicators for working part time and for not working, and their interactions with age 65 or older. One indicator for each age between 26 and 85, and two pooled indicators for ages 25 and below and 86 and above (hollow markers), relative to age 30, the omitted category. The right axis reports the effects as a percentage of the average spending share of one-person households aged 30, the left axis the implied spending share, that average plus the age effect. 95% confidence bands; projection weights; standard errors clustered by household; 298,825 household-years from 58,737 households.

Figure 4: Share of spending on products launched up to five years before by age



Age effects β of equation (1), estimated with all its terms as in Figure 2, for three outcomes: the share of spending on products launched in year t (as in Figure 2), on products launched in the last three years, aged 0 to 2, and on products launched in the last six years, aged 0 to 5. For each outcome, the exposure E_{it} refers to the same products, and the sample keeps the years in which all of them can be observed, 2005–2024, 2007–2024 and 2010–2024, since the first launches are dated 2005. Each line is a percentage of its outcome's average at age 30, which is 0.352%, 1.483% and 3.088% of spending, respectively; hollow markers are the pooled indicators for ages 25 and below and 86 and above. 95% confidence bands; projection weights; standard errors clustered by household.

independent of age. This is the perpetual-youth structure of Yaari (1965) and Blanchard (1985): households are perennially young since their death hazard does not increase with age. The expected lifetime of a household is $1/\mu$ at any time. Households that die are replaced and the inflow of newborns offsets the outflow of deaths, so that the total population is constant and equal to $N > 0$, with $i \in [0, N]$.

Let $a_i^h \geq 0$ denote household i 's age. With a constant death hazard and a stationary population, the cross-sectional distribution of household ages is exponential with rate μ ,

$$a_i^h \sim \text{Exp}(\mu), \quad (2)$$

so that the measure of households whose age is less than or equal to a^h is $N(1 - e^{-\mu a^h})$.

Population ageing is modelled as a fall in the mortality hazard μ , holding the population size N fixed by letting natality adjust one-for-one. A lower μ raises expected lifetime $1/\mu$ and shifts the age distribution (2) towards older households, while leaving the size of the economy unchanged; this isolates the change in age composition from scale effects that would arise from a change in population.²⁶

There is a unit measure of products indexed by $j \in [0, 1]$. A product is described by its quality and by how familiar each consumer is with it. Formally, product j is described by

$$(q_j, [f_{ij}]_{i \in [0, N]}),$$

where $q_j > 0$ is its *quality* and $f_{ij} > 0$ the *familiarity* of consumer i with product j .

Products improve over time through innovation. When innovation happens on product j , it is replaced by a new vintage and

1. the quality of the product increases from q_j to $(1 + g_q)q_j$, where $g_q > 0$;
2. the familiarity of all consumers resets to $f_0 > 0$.

In equilibrium, innovation happens on each product line at the Poisson rate $\tau > 0$, common to all lines. The rate can change over time, and it is constant on a balanced growth path. Let $a_j^p \geq 0$ denote the age of a product j , i.e. the time since the product last underwent innovation.

Households grow familiar with the products they consume, in the tradition of customer-base and deep-habits models (Ravn et al., 2006). A newborn begins her life familiar with all

²⁶Keeping the flow of births fixed, a decrease in μ would increase population N . The Schumpeterian model of growth developed in this paper features scale effects and equilibrium growth rates are increasing in N . Indeed, what matters in terms of scale effects is the total labour force which might decline with ageing for reasons not modelled in the present paper, such as retirement or a concurrent decline in population. A workforce contraction can in turn spur automation and increase productivity (Acemoglu and Restrepo, 2022; Acemoglu et al., 2026). The present paper abstracts from this channel and focuses on new product adoption.

products at the baseline level f_0 . While consumer i is alive and consumes product j ,²⁷ familiarity of i with j grows at the constant rate $g_f > 0$, so that

$$\dot{f}_{ij} = g_f f_{ij}.$$

Hence, familiarity f_{ij} is reset to f_0 by either of two events: (i) consumer i 's death, after which the replacing newborn starts from f_0 ; or (ii) innovation on line j , after which all consumers meet the new vintage at f_0 .

At any point in time, the familiarity of consumer i with product j has therefore been growing at rate g_f since the later of the consumer's birth and the product's last innovation, that is, for $\min\{a_i^h, a_j^p\}$ years, so

$$f_{ij} = f_0 e^{g_f \min\{a_i^h, a_j^p\}}. \quad (3)$$

The model thus features heterogeneity both on the product side, through quality, and on the household side, through age. A key feature that helps make the model tractable is independence between these two types of heterogeneity. The death process and the creative-destruction process are independent of one another, so a_i^h and a_j^p are independent, and household age is independent of everything on the product side. In addition, by construction product age and quality are independent when a product is introduced. However, the quality of a product that has been available on the market for some time is not independent of its age: over the spell its quality is frozen while the rest of the economy keeps innovating, so an older vintage is, in expectation, of lower quality relative to the aggregate.

3.2 Households

Households evaluate consumption streams by expected discounted log utility. For household i born at time t , consuming an amount C_{is} would give a utility flow at time $s \geq t$ equal to $\log C_{is}$. However, the household rationally anticipates that it might not be alive at time s to consume, and it also discounts future consumption relative to present consumption at rate $\rho > 0$. The total discount rate is thus $\rho + \mu$, where μ is the constant death hazard. A consumption stream $[C_{is}]_{s \geq t}$ then has a present discounted value of²⁸

$$\int_t^\infty e^{-(\rho+\mu)(s-t)} \log C_{is} ds. \quad (4)$$

²⁷In equilibrium, every consumer buys every product on the market, although in different quantities.

²⁸Since the utility flow is logarithmic in consumption, as long as consumption grows at most exponentially the presence of a positive discount rate suffices to ensure that the present value of utility is bounded. Generalising to CRRA utility with inverse EIS γ would require assuming the standard restriction $\rho + \mu > (1 - \gamma)\dot{C}_i / C_i$ to ensure bounded utility.

In the remainder of the paper, I drop time indices when this does not cause confusion.

Consumption is a CES aggregate with elasticity $\eta > 1$ over the unit measure of goods, with quality and familiarity entering as standard utility-enhancing weights. Abstracting from time indices, C_i is

$$C_i = \left(\int_0^1 (q_j f_{ij})^{\frac{1}{\eta}} c_{ij}^{\frac{\eta-1}{\eta}} dj \right)^{\frac{\eta}{\eta-1}}. \quad (5)$$

Households supply one unit of labour inelastically at the wage W and save in a risk-free asset A_i . Annuity markets are perfect, so that the risk-free asset earns the actuarially fair return $r + \mu$ (Yaari, 1965; Blanchard, 1985). All firms are owned by a mutual fund that holds the entire equity of the economy, collects aggregate profits Π_t^{agg} , and rebates them as a lump sum to every living household, $\Pi_t = \Pi_t^{\text{agg}} / N$ per household on average. Income can vary with age. A household of age a^h earns the labour income $W(a^h)$ and receives the rebate $\Pi(a^h)$,

$$W(a^h) = \phi(a^h) W, \quad \Pi(a^h) = \phi(a^h) \Pi, \quad \phi(a^h) = \frac{\mu - \phi}{\mu} e^{\phi a^h}, \quad \phi < \mu, \quad (6)$$

where W and Π are the wage and the rebate of the average household. Income rises with age when $\phi > 0$ and falls when $\phi < 0$. The profile $\phi(a^h)$ only redistributes the aggregate wage bill and rebate across ages, since $\int_0^\infty \phi(a^h) \mu e^{-\mu a^h} da^h = 1$, and at $\phi = 0$ every household receives the same income. Below I show that ϕ is also the rate at which expenditure rises with age across households. The profile lets the model match how expenditure is distributed across ages: what matters for the sales of a product is how much households of each age spend, not only how many of them there are. Individual financial wealth therefore evolves according to

$$\dot{A}_i = (r + \mu) A_i + W(a_i^h) + \Pi(a_i^h) - E_i, \quad E_i \equiv \int_0^1 p_j c_{ij} dj, \quad (7)$$

where E_i is the expenditure of household i .

Holding the economy's equity in a fund that rebates dividends to all living households equalises non-human wealth across cohorts and so removes the wealth-replacement channel of Blanchard (1985), leaving demographics to act on the economy only through familiarity.²⁹

Given expenditure E_i , cost minimisation over (5) delivers the demand for each good and the household's ideal price index,

$$c_{ij} = \left(\frac{p_j}{P_i} \right)^{-\eta} q_j f_{ij} C_i = p_j^{-\eta} q_j f_{ij} P_i^{\eta-1} E_i, \quad P_i^{1-\eta} = \int_0^1 q_j f_{ij} p_j^{1-\eta} dj, \quad (8)$$

²⁹If instead equities are traded between households, under perfect annuities newborns enter with zero financial wealth and must accumulate equity at a positive rate, and the shorter the lifespan the faster newborns need to accumulate. The rate of return then depends directly on the mortality rate, creating a second channel through which ageing affects growth. I abstract from this channel in the remainder of the paper.

with $E_i = P_i C_i$. Quality and familiarity enter symmetrically as demand shifters: at given prices a household spends more on goods it knows better or that are of higher quality. In a standard model without familiarity, the price index would be common across households and equal to the baseline $\bar{P}^{1-\eta} = \int_0^1 q_j p_j^{1-\eta} dj$. For a newborn household that is familiar with all products at the reset value f_0 , the price index factorises into familiarity and the baseline

$$P_i^{1-\eta} = f_0 \bar{P}^{1-\eta} . \quad (9)$$

Because f_{ij} rises with the household's age, the ideal price index P_i depends on age, with older consumers facing a lower ideal price index. However, no such factorisation is available for older consumers: they are more familiar with older products, whose quality has been frozen since their last innovation while new vintages kept raising the aggregate, so that on average they are of lower quality. The dependence of the price index on age is analytically derived in Section 3.7.

With logarithmic utility, solving for the optimal expenditure path $\{E_i\}$ subject to (7) delivers the standard solution. The costate λ_i of the budget constraint satisfies $\dot{\lambda}_i / \lambda_i = (\rho + \mu) - (r + \mu) = \rho - r$ and $\lambda_i = 1/E_i$, so that the Euler equation is

$$\frac{\dot{E}_i}{E_i} = r - \rho, \quad (10)$$

where the death hazard cancels between the annuity return $r + \mu$ and the effective discount $\rho + \mu$.³⁰

Familiarity does not appear in the Euler equation. Since $E_i = P_i C_i$, the utility flow is the log of expenditure minus the log of the household's ideal price index. Section 3.7 shows that, on a balanced growth path with a common price and the normalisation $\bar{P} = 1$, the price index is $P_i^{1-\eta} = \mathcal{F}^h(a_i^h)$ (25), where $\mathcal{F}^h(a_i^h)$ is the average familiarity of a household of age a_i^h with the products on the market (24). The utility flow is then

$$\log C_i = \log E_i + \frac{1}{\eta - 1} \log \mathcal{F}^h(a_i^h) . \quad (11)$$

Average familiarity rises with age, so each unit of expenditure buys more consumption as a household grows older. One may then ask why households do not save to spend more later in life, when they are more familiar with the products they buy. The answer lies in logarithmic utility. In (11), familiarity enters additively, and it does not depend on the household's choices,

³⁰Absent annuities, a decline in the death hazard would make households discount the future less, similarly to an increase in patience. Given any growth rate, the rate of return would decrease, stimulating entry by increasing the value of a firm in the economy.

since familiarity grows at the rate g_f whatever the quantity consumed. The marginal utility of expenditure is therefore $1/E_i$ at every age, whatever the household's familiarity. Rising familiarity acts like a higher real return on saving, and with logarithmic utility its substitution effect, which favours spending later, exactly offsets its income effect, which favours spending now. Familiarity changes how a household divides its expenditure across products, not how it divides it over time. The logarithmic specification is therefore important: with any other elasticity of intertemporal substitution, familiarity would also shape saving and the age profile of expenditure.³¹

Under the CES aggregator, familiarity with a product (and hence age) increases the amount of that product consumed, but every household consumes every product. The model therefore has no extensive margin of new product adoption, only an intensive one: older households buy less of a new product, rather than not buying it at all. This keeps the model tractable and easy to aggregate. Its counterpart in the data is the share of spending on new products of Section 2.2, which reflects both margins.

Real consumption of household i then grows at $\dot{C}_i/C_i = (r - \rho) - \dot{P}_i/P_i$, where the second term captures the decline in the household's ideal price index as familiarity accumulates. Integrating (7) forward under the standard no-Ponzi condition gives the consumption function under log utility

$$E_{it} = (\rho + \mu)[A_{it} + H_{it}], \quad H_{it} = \int_t^\infty [W_s(a_i^h + s - t) + \Pi_s(a_i^h + s - t)] e^{-(r+\mu)(s-t)} ds, \quad (12)$$

with H_{it} denoting human wealth. Under perpetual youth and the mutual fund, human wealth depends on age only through the income profile: on a path on which W and Π grow at the rate g , a household's income grows at the rate $\phi + g$ over its life, and $H_{it} = [W_t(a_i^h) + \Pi_t(a_i^h)]/(r + \mu - \phi - g)$. A newborn and an old household differ in human wealth only because they differ in current income.

³¹With CRRA utility $C_i^{1-\gamma}/(1-\gamma)$, the marginal utility of expenditure is $E_i^{-\gamma} \mathcal{F}^h(a_i^h)^{\frac{1-\gamma}{\eta-1}}$, and the Euler equation becomes

$$\frac{\dot{E}_i}{E_i} = \frac{r - \rho}{\gamma} + \frac{1 - \gamma}{\gamma(\eta - 1)} \frac{\dot{\mathcal{F}}^h(a_i^h)}{\mathcal{F}^h(a_i^h)}.$$

The second term is new. If the elasticity of intertemporal substitution $1/\gamma$ is above one, households save to spend more when they are older and more familiar with products; if it is below one, they spend more while young. Either way, expenditure growth depends on age through the growth rate of average familiarity, which equals g_f at birth and declines with age. Households of different ages would then hold different amounts of the bond, expenditure would no longer follow the income profile (6), and the interest rate would no longer be given by (13). Logarithmic utility is also why the market's static allocation of goods across households coincides with the planner's in Section 8 (Proposition 8): since familiarity shifts the log of consumption additively, it does not change the marginal utility of goods, and the planner has no reason to give more goods to the older households, who are more familiar with the products on the market.

At equilibrium, the risk-free bond is in zero net supply, and every household holds zero savings. A household with no assets spends its income, $E_{it} = W_t(a_i^h) + \Pi_t(a_i^h)$, which grows at the rate $\phi + g$ over its life. By the Euler equation (10), this is optimal if and only if

$$r = \rho + \phi + g, \quad (13)$$

and the consumption function (12) then gives the same expenditure. Expenditure therefore follows the income profile (6), $E_i = \phi(a_i^h)E$, where E is expenditure per household: across households, it rises with age at the rate ϕ . Households whose expenditure rises over life behave as if they were less patient, since the gap between the return and growth is $r - g = \rho + \phi$. At $\phi = 0$ every household spends the same, $E_i = E$, and $r - g = \rho$.

3.3 Demand aggregation and the familiarity index

The original feature of the model presented so far is that consumer-specific familiarity with products sold on the market affects the quantity each consumer purchases. This, in turn, implies that the demand curve faced by the seller of a product depends on familiarity. Familiarity affects how much of a specific product is sold in two ways: *i*) the more familiar a product is to consumers, the more it will sell at any given price; *ii*) the more familiar consumers are with *all other goods* sold in the market, the less the good will sell. Hence, even though the consumption aggregator is CES, a crucial and non-trivial step in solving the model is to aggregate the demands of consumers $i \in [0, N]$ for good j into the aggregate demand that j 's producer faces.

To start, I integrate individual demands (8) over $[0, N]$, using $E_i = \phi(a_i^h)E$, to get the total demand from consumers for good j ,

$$c_j = \int_0^N c_{ij} \, di = \int_0^N p_j^{-\eta} q_j f_{ij} P_i^{\eta-1} \phi(a_i^h) E \, di = p_j^{-\eta} q_j E \int_0^N f_{ij} P_i^{\eta-1} \phi(a_i^h) \, di.$$

I now define the *familiarity index* $\mathcal{F}$ as

$$\mathcal{F}(j) = \frac{1}{N} \int_0^N f_{ij} \left(\frac{\bar{P}}{P_i} \right)^{1-\eta} \phi(a_i^h) \, di.$$

The specific functional form of the index is justified by the following normalisation: shut down familiarity growth, so that for all consumers and goods $f_{ij} = f_0$, then $P_i^{1-\eta} = f_0 \bar{P}^{1-\eta}$ (9) and $\mathcal{F}(j) = 1$, since $\phi(a^h)$ integrates to one over the population. Each household enters the index with the weight of its expenditure.

Under this definition I can write the consumption demand for good j as

$$c_j = \left(\frac{p_j}{\bar{p}} \right)^{-\eta} q_j \mathcal{F}(j) \frac{NE}{\bar{p}}.$$

Substituting out the expressions for the individual and baseline ideal price indices shows clearly how familiarity affects demand in two directions,

$$\mathcal{F}(j) = \frac{1}{N} \int_0^N \underbrace{f_{ij}}_+ \frac{\int_0^1 q_{j'} p_{j'}^{1-\eta} dj'}{\int_0^1 \underbrace{f_{ij'}}_+ q_{j'} p_{j'}^{1-\eta} dj'} \phi(a_i^h) di, \quad (14)$$

own familiarity f_{ij} increases sales, while familiarity of competitors $f_{ij'}$ decreases them. The familiarity index also has simple properties: it is homogeneous of degree 1 in the distribution of own familiarity $[f_{ij}]_{i \in [0, N]}$, of degree -1 in the familiarity with other products $[f_{ij'}]_{i \in [0, N], j' \neq j}$ and of degree 0 in other prices $[p_{j'}]_{j' \neq j}$ and qualities $[q_{j'}]_{j' \neq j}$. In addition, since the familiarity index is homogeneous of degree 1 in own familiarity and -1 in other products' familiarity, it follows that it is homogeneous of degree 0 in the full familiarity schedule $[f_{ij}]_{i \in [0, N], j \in [0, 1]}$. This implies that a rescaling of familiarity has no effect on allocations.

So far, I have only aggregated demand for household consumption. However, goods produced in the economy can be used for two purposes: consumption and research, performed by potential entrants on the firm side. For tractability, I assume that the quantity of individual goods used in research scales proportionally with the quantity used by households,³² so that I can write the demand for good j as

$$y_j = \left(\frac{p_j}{\bar{p}} \right)^{-\eta} q_j \mathcal{F}(j) \frac{Y}{\bar{p}}, \quad (15)$$

where Y denotes total expenditure (household consumption plus research spending).

Given this demand, the price-setting problem of incumbent firms is simple: each firm takes its familiarity index as given and faces an isoelastic demand, so it sets its price at a constant markup above marginal cost. Since the only input of production is labour and the wage is common across firms, all firms choose the same price. The problem of the firm is formally described and solved in the following section of the paper. An explicit expression for the familiarity index

³²The formal assumption that ensures this is that research requires a good X which is determined according to

$$X = \left(\int_0^1 (q_j \mathcal{F}(j))^{\frac{1}{\eta}} x_j^{\frac{\eta-1}{\eta}} dj \right)^{\frac{\eta}{\eta-1}},$$

where x_j is the quantity of good j used in research.

requires this common price and the distributions of household and product ages on a balanced growth path, so I derive it in Section 3.7, once the equilibrium is defined.

3.4 Incumbents

Each good $j \in [0, 1]$ is produced by one incumbent firm. The firm is a monopolist that maximises profits choosing prices and quantities under the demand schedule. While in principle the good is characterised by its quality and familiarity distribution with consumers, $(q_j, [f_{ij}]_{i \in [0, N]})$, the quantities relevant to the firm problem are its quality q_j and its familiarity index $\mathcal{F}(j)$, which summarises that distribution and which the firm takes as given.

Monopolistic firms produce one unit of good using one unit of labour,

$$y_j = l_j ,$$

pay the wage W (so that the marginal cost is constant and equal for all producers), and face the demand (15). I take the baseline price index as the numeraire, $\bar{P} = 1$, so that the demand for good j is $y_j = p_j^{-\eta} q_j \mathcal{F}(j) Y$.

The problem of monopolist j is therefore

$$\pi_j = \max_{p, y, l} py - Wl , \quad \text{s.t.} \quad y = p^{-\eta} q_j \mathcal{F}(j) Y , \quad y = l .$$

The firm chooses the price to be a constant markup over the marginal cost

$$p = \frac{\eta}{\eta - 1} W , \tag{16}$$

which implies that profits are

$$\pi_j = q_j \mathcal{F}(j) \Pi , \quad \Pi \equiv \frac{1}{\eta} \left(\frac{\eta}{\eta - 1} \right)^{1-\eta} W^{1-\eta} Y ,$$

so that profits are a function of quality, the familiarity index and aggregate quantities. In principle, an incumbent could be constrained in its price setting by the presence of previous producers along the same product line. To keep the model tractable, I adopt the following assumption.

Assumption 1 (Monopoly pricing). Each monopolistic producer is unconstrained by previous producers along the same product line and can charge the monopoly price (16).

Limit pricing, which is the equilibrium alternative to monopoly pricing in standard quality ladder models, would be particularly difficult to solve in the present setting, since there

is no single price, independent of history and time, that a new producer can charge to completely displace the previous incumbent. This is an implication of having a familiarity distribution of consumers with products that evolves over time: a new product enters the utility of all households with weight $(1 + g_q)q_j f_0$, while the previous product has a consumer-specific weight $q_j f_{ij}$. Hence, at equal prices newborns prefer the new product, while consumers with $\min\{a_i^h, a^{p,old_j}\} > \ln(1 + g_q)/g_f$ prefer the old vintage. As a result, the price that captures full demand of the previous incumbent depends on the age of the previous product, and would also evolve with time as familiarity with the new product increases. In addition, this form of limit pricing would require keeping track of the familiarity distributions of previous products and making assumptions on how familiarity is remembered, increasing the state space.³³

The original part of the firm value lies in how profits change over the life of the product, as consumers become familiar with it. The value of the firm V_j is determined as usual by the HJB equation

$$(r + \tau)V_j - \frac{dV_j}{dt} = \pi_j .$$

While quality q_j is fixed, the familiarity index of the product changes over time as consumers become familiar with it, and this appears in dV_j/dt . Since profits depend on familiarity only through the index, I can also write $V_j = V(q_j, \mathcal{F}(j); W, Y)$, so that

$$(r + \tau)V_j - \dot{\mathcal{F}}(j) \partial_{\mathcal{F}(j)} V_j - \dot{V}_j = \pi_j , \quad (17)$$

excluding familiarity from the dot notation in $\dot{V}_j$. On a balanced growth path, $\mathcal{F}(j) = \mathcal{F}(a_j^p)$ (Section 3.7), and $\dot{\mathcal{F}}(j) = \mathcal{F}'(a_j^p)$, since age is linear in time.

3.5 Entry

A mass of entrants invests in research to create new products and become monopolistic producers upon successful innovation, displacing previous incumbents. A successful entrant carries out product innovation on a line j drawn uniformly on the product space $[0, 1]$ and receives a patent to produce it at the quality $(1 + g_q)q_j$. At the same time, product innovation resets the familiarity of all consumers with product j to the baseline f_0 , so that its familiarity index starts from $\mathcal{F}^0 \equiv \frac{1}{N} \int_0^N f_0 (\bar{P}/P_i)^{1-\eta} \phi(a_i^h) di$, the same for every new product. Entrants choose the innovation flow rate $\tau \geq 0$ with the linear R&D cost $X(\tau) = v\tau\bar{Q}$, where v is a scale parameter and $\bar{Q}$ is the *quality index*,

$$\bar{Q}_t^{\eta-1} \equiv \mathbb{E}[q_{jt}] = \int_0^1 q_{jt} dj .$$

³³An alternative possibility would be to allow old incumbents to survive along with new ones and slowly fade away. Section 8 discusses how this assumption shapes the familiarity-destruction externality.

The exponent $\eta - 1$ is a normalisation which makes this quantity particularly relevant for the analysis: in equilibrium, key aggregates such as total expenditure and the wage scale with $\bar{Q}$. The scaling of the cost by aggregate quality ensures a constant creative destruction rate at equilibrium. The value V^e of being a potential entrant is thus the expected value of being successful in innovating on a product and becoming a monopolistic producer, and is determined by

$$rV^e - \dot{V}^e = \max_{\tau \geq 0} \left\{ \tau \left(\mathbb{E}_j[V((1 + g_q)q_j, \mathcal{F}^0)] - V^e \right) - v\tau\bar{Q} \right\}. \quad (18)$$

At equilibrium, free entry ensures $V^e = 0$; when there is positive entry the equilibrium is such that

$$\mathbb{E}_j[V((1 + g_q)q_j, \mathcal{F}^0)] = v\bar{Q}. \quad (19)$$

3.6 Market clearing and equilibrium definition

Labour market clearing in the economy requires

$$\int_0^1 l_j \, dj = N,$$

and the value of output equal to total spending implies

$$Y = NE + X$$

where E is expenditure per household and X is the research spending.

The risk-free asset is in zero net supply

$$A = 0,$$

and the mutual fund owns all firms and rebates profits to living households as described in Section 3.2.

I next define a balanced growth path equilibrium of the economy.

Definition 1 (Balanced growth path equilibrium). A *balanced growth path* (BGP) consists of household allocations $\{c_{ij}, E_i, A_i\}$ and price indices $\{P_i\}$; incumbent prices, outputs and employment $\{p_j, y_j, l_j\}$ with value function V_t ; a creative-destruction rate τ , research spending X and an entrant value V^e ; aggregates $\{Y_t, W_t, \bar{Q}_t, C_t\}$, an interest rate r and a growth rate g ; and cross-sectional distributions of household ages a^h and product ages a^p such that, at every date t :

- (i) *Households*. $\{c_{ij}, E_i, A_i\}$ solve the problem of Section 3.2, satisfying demand (8), the Euler

equation (10) and the consumption function (12).

- (ii) *Incumbents.* Prices satisfy (16), outputs (15), employment $l_j = y_j$, and V_t solves (17).
- (iii) *Entry.* $V^e = 0$ solves (18).
- (iv) *Market clearing.* Labour, goods and asset markets clear.
- (v) *Stationarity.* The age distribution of households is the stationary exponential (2), and that of products is exponential with rate τ .
- (vi) *Balanced growth.* $\dot{Y}_t/Y_t = \dot{W}_t/W_t = \dot{C}_t/C_t = \dot{\bar{Q}}_t/\bar{Q}_t = g$, constant, while r and τ are constant over time.

3.7 Equilibrium characterisation and growth

Growth of the quality index. On a balanced growth path, the rate of creative destruction pins down the growth of the quality index.

Lemma 1. *Suppose that the rate of creative destruction τ is constant and innovators innovate on random lines, then*

$$g_{\bar{Q}} = \frac{\tau g_q}{\eta - 1}. \quad (20)$$

Proofs not given in the text, starting with this one, are in Appendix C.

Each term in (20) has a simple reading. In a unit of time, a share τ of the product lines receives a new vintage. Since innovators pick lines at random, the lines they improve have the average quality, and each innovation raises the quality of its line by the step g_q . Average quality, $\int_0^1 q_j dj = \bar{Q}^{\eta-1}$, therefore grows at the rate τg_q . The exponent $\eta - 1$ converts quality into consumption. Quality enters demand as a weight: for a household equally familiar with two goods, a good of quality q_j sold at the price p_j is worth as much as a good of quality one sold at $p_j / q_j^{1/(\eta-1)}$, since the two contribute equally to the price index (8). The quality index $\bar{Q}$ is thus the quality-adjusted output of one unit of labour, and output is $Y = N\bar{Q}$ in equilibrium, as shown below. It grows at $1/(\eta - 1)$ times the rate of average quality: the more substitutable goods are, the less consumption a given quality step adds. Section 4.2 uses (20) to recover the quality step from the rates of growth and firm entry.

The familiarity index on a balanced growth path. To characterise the equilibrium defined in the previous section, I begin from the normalisation $\bar{P}^{1-\eta} = \int_0^1 q_j p_j^{1-\eta} dj = 1$ of Section 3.4. This immediately implies that the monopolistic price, common to all incumbents, is $p = \bar{Q}$.

With the common price, (14) simplifies to

$$\mathcal{F}(j) = \frac{1}{N} \int_0^N f_{ij} \frac{\int_0^1 q_{j'} dj'}{\int_0^1 q_{j'} f_{ij'} dj'} \phi(a_i^h) di .$$

The following step is to change variables, from indices i and j to household and product ages, integrating over their distribution. On a balanced growth path, product ages are exponential with rate τ , independent of household ages. Since each household is weighted by its expenditure, the relevant density of household ages is the expenditure-weighted one,

$$\phi(a^h) \mu e^{-\mu a^h} = \tilde{\mu} e^{-\tilde{\mu} a^h} , \quad \tilde{\mu} \equiv \mu - \phi , \quad (21)$$

again exponential, with the rate $\tilde{\mu}$. I call $\tilde{\mu}$ the *spending hazard*: demand behaves as if households died at the rate $\tilde{\mu}$. When expenditure rises with age, $\phi > 0$, spending is concentrated on older households, and demand behaves as if the population were older than it is. Then I can use the law of motion of familiarity to express $f_{ij} = f_0 e^{g_f \min\{a_i^h, a_j^p\}}$. Doing so gives the following expression³⁴

$$\mathcal{F}(a^p) = \int_0^\infty f_0 e^{g_f \min\{a^h, a^p\}} \frac{\mathbb{E}[q_{j't}]}{\int_0^\infty \mathbb{E}[q_{j't} | a_{j'}^p = a^{p'}] f_0 e^{g_f \min\{a^h, a^{p'}\}} \tau e^{-\tau a^{p'}} da^{p'}} \tilde{\mu} e^{-\tilde{\mu} a^h} da^h .$$

By the law of iterated expectations, $\mathbb{E}[q_{j't}] = \int_0^\infty \mathbb{E}[q_{j't} | a_{j'}^p = a^{p'}] \tau e^{-\tau a^{p'}} da^{p'}$. However, I cannot factorise the second term out of the denominator in the expression above since, as noted above, the age and the quality of a product are not independent, in a sense made precise now. Define the *conditional quality index* $Q_t(a^p)$ by

$$Q_t(a^p)^{\eta-1} \equiv \mathbb{E}[q_{jt} | a_{jt}^p = a^p] ,$$

which captures the average quality at time t of a line of age a^p in units of the quality index $\bar{Q}_t$ of Section 3.5. The law of iterated expectations implies $\bar{Q}_t^{\eta-1} = \int_0^\infty Q_t(a^p)^{\eta-1} \tau e^{-\tau a^p} da^p$. The following lemma characterises how the conditional quality index evolves with age, formalising how quality of a product falls behind as it ages. Importantly, the result hinges on the fact that a successful entrant improves a random product line by the step g_q .

³⁴Notice that the $1/N$ term in front of the familiarity index disappears when changing variables, since $di = N \mu e^{-\mu a^h} da^h$, and by (21) the weights become $\tilde{\mu} e^{-\tilde{\mu} a^h}$.

Lemma 2. *Along a balanced growth path on which $\bar{Q}_t$ grows at the constant rate $g_{\bar{Q}}$,*

$$Q_t(a^p)^{\eta-1} = (1 + g_q) e^{-(\eta-1)g_{\bar{Q}}a^p} \bar{Q}_t^{\eta-1}.$$

At $a^p = 0$, innovation happening on a random line delivers the conditional mean $(1 + g_q)\bar{Q}_t^{\eta-1}$: a new vintage is, on average, one quality step ahead of the average. As the vintage ages, its own quality is frozen while the aggregate keeps growing, so its relative quality decays at the rate $(\eta - 1)g_{\bar{Q}}$. Using this result, I can rewrite the familiarity index as

$$\mathcal{F}(a^p) = \int_0^\infty \frac{e^{g_f \min\{a^h, a^p\}}}{\int_0^\infty e^{g_f \min\{a^h, a^{p'}\}} (1 + g_q) \tau e^{-(\tau + (\eta-1)g_{\bar{Q}})a^{p'}} da^{p'}} \tilde{\mu} e^{-\tilde{\mu}a^h} da^h,$$

where f_0 has also been simplified. The expression in the denominator says that since older competing lines have lower quality, their negative demand effect on a good of age a^p is discounted. In fact, the density of unit mass $\tau e^{-\tau a^p}$ over $a^p \in [0, \infty)$ gets transformed into another density of the same mass, since at equilibrium $(1 + g_q)\tau e^{-(\tau + (\eta-1)g_{\bar{Q}})a^p}$ integrates to one. This is a corollary of Lemma 1, which implies that $\tau + (\eta - 1)g_{\bar{Q}} = (1 + g_q)\tau$, so that I can express

$$\mathcal{F}(a^p) = \int_0^\infty \frac{e^{g_f \min\{a^h, a^p\}}}{\int_0^\infty e^{g_f \min\{a^h, a^{p'}\}} (1 + g_q) \tau e^{-(1+g_q)\tau a^{p'}} da^{p'}} \tilde{\mu} e^{-\tilde{\mu}a^h} da^h. \quad (22)$$

Relative to the raw age distribution $\tau e^{-\tau a^p}$, the density $(1 + g_q)\tau e^{-(1+g_q)\tau a^p}$ is exponential with the faster rate $(1 + g_q)\tau > \tau$ and “stochastically younger”. For the purpose of demand, the familiarity stock needs to be quality-weighted, ending up effectively younger than its age distribution. It also follows that

$$\int_0^\infty \mathcal{F}(a^p) (1 + g_q) \tau e^{-(1+g_q)\tau a^p} da^p = 1. \quad (23)$$

Given the parameters, $\mathcal{F}$ therefore reallocates demand across products of different ages, while total demand, weighted with the quality-adjusted age density, stays constant.³⁵ The important distinction is that the quality-weighted density of product ages depends on τ , but not on the spending hazard $\tilde{\mu}$, while $\mathcal{F}$ depends on $\tilde{\mu}$, and on τ only through that density. In practice, given τ , $\mathcal{F}$ defines a one-parameter family of demand densities in $\tilde{\mu}$. Ageing and the expenditure profile act on demand only through $\tilde{\mu} = \mu - \phi$.

The following proposition describes the main properties of the familiarity index $\mathcal{F}$.

Proposition 1. *For all parameter values, the familiarity index satisfies:*

³⁵Formally, $\mathcal{F}$ is the Radon-Nikodym derivative of the resulting density with respect to the quality-weighted density of product ages.

- (i) $\mathcal{F}$ is finite and strictly increasing, with $\mathcal{F}'(0) = g_f \mathcal{F}(0)$. Its growth rate satisfies $0 < \mathcal{F}'(a^p)/\mathcal{F}(a^p) \leq g_f$, with equality only at $a^p = 0$;
- (ii) $\mathcal{F}(0) < 1 < \mathcal{F}(\infty) \equiv \lim_{a^p \rightarrow \infty} \mathcal{F}(a^p) \leq \infty$, with $\mathcal{F}(\infty) < \infty$ if and only if $\min\{g_f, (1+g_q)\tau\} < \tilde{\mu}$;
- (iii) $\partial_{\tilde{\mu}} \mathcal{F}(0) > 0$ and there exists a unique $a^{p*} > 0$ such that $\partial_{\tilde{\mu}} \mathcal{F}(a^p)$ is positive below a^{p*} and negative above it: ageing lowers familiarity of young products, and raises it for older ones, and so does a steeper expenditure profile;
- (iv) $\partial_{\tau} \mathcal{F}(a^p) > 0$ for every $a^p \geq 0$: a faster replacement rate raises the familiarity index at every age level.

The proposition requires no restriction on the parameters. By (i), $\mathcal{F}(a^p) \leq e^{g_f a^p} \mathcal{F}(0)$, so the index is finite at every product age, but it grows without bound as the product ages when $\min\{g_f, (1+g_q)\tau\} \geq \tilde{\mu}$.

Finally, I can describe how the ideal price index of a household evolves as the household becomes older. Define the average familiarity of a household of age a^h ,

$$\mathcal{F}^h(a^h) \equiv \int_0^\infty f_0 e^{g_f \min\{a^h, a^p\}} (1+g_q)\tau e^{-(1+g_q)\tau a^p} da^p, \quad (24)$$

the quality-weighted mean of its familiarity schedule (3) over products. With the common price,

$$P(a^h)^{1-\eta} = \int_0^1 q_j f_{ij} p_j^{1-\eta} dj = p^{1-\eta} \bar{Q}^{\eta-1} \mathcal{F}^h(a^h).$$

Under the normalisation $\bar{P}^{1-\eta} = \bar{Q}^{\eta-1} p^{1-\eta} = 1$, this becomes

$$P(a^h)^{1-\eta} = \mathcal{F}^h(a^h). \quad (25)$$

Then, if $g_f < (1+g_q)\tau$

$$\mathcal{F}^h(a^h) = f_0 \frac{(1+g_q)\tau - g_f e^{-((1+g_q)\tau - g_f)a^h}}{(1+g_q)\tau - g_f},$$

so that the ideal price index of a household decreases with age and converges to a positive constant. If instead $g_f > (1+g_q)\tau$, the price index converges to zero, asymptotically at exponential rate $(g_f - (1+g_q)\tau)/(\eta-1)$. If equality holds, the price index decays to zero but less than exponentially.

On a balanced growth path, the familiarity index of a product therefore depends on the

product only through its age, $\mathcal{F}(j) = \mathcal{F}(a_j^p)$, and the demand (15) for good j is

$$y_j = p_j^{-\eta} q_j \mathcal{F}(a_j^p) Y. \quad (26)$$

Aggregates and the value of firms. Labour market clearing implies

$$N = \int_0^1 l_j dj = \int_0^1 p^{-\eta} q_j \mathcal{F}(j) Y dj = p^{-\eta} Y \bar{Q}^{\eta-1}$$

so that total output is

$$Y = N \bar{Q}.$$

The monopolist price implies also that

$$W = \frac{\eta-1}{\eta} \bar{Q} = \frac{\eta-1}{\eta} \frac{Y}{N}, \quad \Pi = \frac{1}{\eta} N \bar{Q} = \frac{1}{\eta} Y.$$

The labour share of the model is $(\eta-1)/\eta$ with remaining output being distributed as profits.

With wages and output expressed in terms of aggregate quality, the profits of firm j are

$$\pi_j = \frac{1}{\eta} q_j \mathcal{F}(a_j^p) N \bar{Q}^{2-\eta}. \quad (27)$$

Profits grow with $\bar{Q}^{2-\eta}$ and are discounted at the rate $r + \tau$. By (13) and Lemma 1, the net rate is $r + \tau - (2-\eta)g = \tilde{\rho} + (1+g_q)\tau$, where $\tilde{\rho} \equiv \rho + \phi = r - g$ is the gap between the return and growth. I can then use guess-and-verify to study the value of the firm, as in the following proposition.

Proposition 2. *Along the balanced growth path the value function of the firm is linear in profits and can be written as*

$$V(q_j, a_j^p) = \kappa(a_j^p) \pi_j, \quad (28)$$

where κ is the unique solution, under the transversality condition $\lim_{u \rightarrow \infty} e^{-(\tilde{\rho} + (1+g_q)\tau)u} \kappa(a^p + u) \mathcal{F}(a^p + u) = 0$, of

$$(\tilde{\rho} + (1+g_q)\tau) \kappa(a^p) - \kappa'(a^p) - \kappa(a^p) \frac{\mathcal{F}'(a^p)}{\mathcal{F}(a^p)} = 1, \quad (29)$$

namely the exact present value of the line's familiarity path relative to its current level,

$$\kappa(a^p) = \int_0^\infty e^{-(\tilde{\rho} + (1+g_q)\tau)u} \frac{\mathcal{F}(a^p + u)}{\mathcal{F}(a^p)} du. \quad (30)$$

Moreover, under $g_f < \tau(1+g_q)$:

(i) κ is differentiable, with

$$\kappa'(a^p) = \int_0^\infty e^{-(\tilde{\rho} + (1+g_q)\tau)u} \frac{\mathcal{F}(a^p + u)}{\mathcal{F}(a^p)} \left[\frac{\mathcal{F}'(a^p + u)}{\mathcal{F}(a^p + u)} - \frac{\mathcal{F}'(a^p)}{\mathcal{F}(a^p)} \right] du; \quad (31)$$

(ii) $\kappa'(0) < 0$.

The main result in the proposition above is that, unlike in a standard quality ladder model, or a model in which firms accumulate product lines as in Klette and Kortum (2004), the value of a firm, while linear in profits, features an age-varying component which captures sales growth through familiarity. This is a natural feature of asset prices: for a firm of age zero, the value of its stock is partly due to current profits, and partly a promise of future profits. As profits start to rise, the forward-looking component starts to fall ($\kappa'(0) < 0$). The value multiple $\kappa = V/\pi$ is a price-earnings ratio.

Free entry, existence and uniqueness of BGP. I can now combine free entry with the expression for the firm value. Using (28) in (19), with $\mathcal{F}^0 = \mathcal{F}(0)$ on a balanced growth path, I find

$$\kappa(0)(1 + g_q) \int_0^1 q_j dj \mathcal{F}(0) N \bar{Q}^{2-\eta} = \nu \bar{Q}$$

which implies

$$\kappa(0) \mathcal{F}(0) = \frac{\bar{\nu}}{(1 + g_q)}, \quad \bar{\nu} \equiv \frac{\eta \nu}{N}. \quad (32)$$

Free entry equates the entrant's value to the composite entry cost $\bar{\nu}$. The parameters ν and N enter the equilibrium only through the composite; a larger population lowers $\bar{\nu}$ and stimulates entry. To evaluate the left-hand side of (32), set $a^p = 0$ in the present value (30): $\kappa(0)$ is the familiarity path of a new product relative to its initial level, discounted at the net rate $\tilde{\rho} + (1 + g_q)\tau$. Multiplying by $\mathcal{F}(0)$ gives the discounted familiarity path itself, a Laplace transform,

$$\kappa(0) \mathcal{F}(0) = \int_0^\infty e^{-(\tilde{\rho} + (1+g_q)\tau)u} \mathcal{F}(u) du = \widehat{\mathcal{F}}(\tilde{\rho} + (1 + g_q)\tau),$$

where $\widehat{\cdot}$ denotes the Laplace transform of a function, so that I can write the free entry condition

$$\widehat{\mathcal{F}}(\tilde{\rho} + (1 + g_q)\tau; \tilde{\mu}, \tau) = \frac{\bar{\nu}}{1 + g_q}. \quad (33)$$

The value of entering, on the left-hand side, adds up the familiarity index of a new product along its life, that is, how much the product sells at each age relative to an economy in which households never become familiar with products. Sales at later ages are discounted at the rate

$\tilde{\rho} + (1 + g_q)\tau$, since the return exceeds growth by $\tilde{\rho}$, the product is replaced at the rate τ , and its share of spending falls at the rate $(\eta - 1)g = g_q\tau$ as other lines improve, so the sales a product gains once households have become familiar with it count for less than its sales at launch. This is the equation that determines the creative destruction rate and in which the notation makes clear that $\widehat{\mathcal{F}}$ depends on τ both via $\mathcal{F}$ and via the weight of the transform. Everything else on the balanced growth path follows from solving this equation.

The entry value is finite even when $\mathcal{F}$ is unbounded. By the normalisation (23), its integral discounted at $(1 + g_q)\tau$ equals $1/[(1 + g_q)\tau]$, and the additional discount $\tilde{\rho} = \rho + \phi > 0$ makes the entry value smaller, which requires $\phi > -\rho$. Proposition 3 below gives a sharper bound and characterises when entry is positive.

A balanced growth path with positive consumption must also satisfy one more condition. Goods market clearing gives

$$NE = Y - X(\tau) = \left(1 - \frac{\bar{v}\tau}{\eta}\right) Y \quad (34)$$

so that positive consumption requires the feasibility condition $\bar{v}\tau < \eta$.³⁶

The solution of the model hinges on finding a τ that solves (33). The next proposition establishes when solutions exist.

Proposition 3. *For all parameter values with $-\rho < \phi < \mu$ the entry value $\widehat{\mathcal{F}}(\tilde{\rho} + (1 + g_q)\tau; \tilde{\mu}, \tau)$ is finite, strictly positive and continuous in τ on $(0, \infty)$, and satisfies*

$$0 < \widehat{\mathcal{F}}(\tilde{\rho} + (1 + g_q)\tau; \tilde{\mu}, \tau) < \frac{1}{\tilde{\rho} + (1 + g_q)\tau}, \quad (35)$$

with boundary limits

$$\lim_{\tau \rightarrow 0^+} \widehat{\mathcal{F}}(\tilde{\rho} + (1 + g_q)\tau; \tilde{\mu}, \tau) = \frac{\tilde{\mu}(\tilde{\mu} + \tilde{\rho} + g_f)}{\tilde{\rho}(\tilde{\mu} + \tilde{\rho})(\tilde{\mu} + g_f)}, \quad \lim_{\tau \rightarrow \infty} \widehat{\mathcal{F}}(\tilde{\rho} + (1 + g_q)\tau; \tilde{\mu}, \tau) = 0. \quad (36)$$

Hence $\bar{V} \equiv \sup_{\tau > 0} \widehat{\mathcal{F}}(\tilde{\rho} + (1 + g_q)\tau; \tilde{\mu}, \tau)$ is finite. If $\bar{v} < (1 + g_q)\bar{V}$, the entry condition (33) has at least one solution $\tau \in (0, \infty)$; every solution satisfies $\bar{v}\tau < 1 < \eta$, so that feasibility holds and each solution is a BGP with positive entry. If $\bar{v} > (1 + g_q)\bar{V}$, no BGP with positive entry exists.

By analogy with other quality ladder models, one might conjecture $\widehat{\mathcal{F}}(\tilde{\rho} + (1 + g_q)\tau; \tilde{\mu}, \tau)$ to be monotone decreasing in τ , since a higher rate of creative destruction reduces the value of a firm by shortening expected lifespan. However, in the present model a higher entry rate has a second effect, since it increases demand by lowering the familiarity stock with which entrants

³⁶This condition never restricts the set of balanced growth paths: Proposition 3 below shows that it is satisfied automatically at every solution of the entry condition (33), for any parameter values.

compete. The two effects can be identified by defining the general transform $\widehat{\mathcal{F}}(\sigma; \tilde{\mu}, \tau)$ and deriving the left-hand side of (33),

$$\begin{aligned} \frac{d}{d\tau} \widehat{\mathcal{F}}(\tilde{\rho} + (1 + g_q)\tau; \tilde{\mu}, \tau) = & \underbrace{\int_0^\infty e^{-(\tilde{\rho} + (1 + g_q)\tau)u} \partial_\tau \mathcal{F}(u) du}_{\partial_\tau \widehat{\mathcal{F}} > 0} \\ & - (1 + g_q) \underbrace{\int_0^\infty u e^{-(\tilde{\rho} + (1 + g_q)\tau)u} \mathcal{F}(u) du}_{\partial_\sigma \widehat{\mathcal{F}} < 0}, \end{aligned} \quad (37)$$

Whether the equilibrium solution is unique then becomes a question of whether one of the two effects dominates on the relevant domain for τ . For sufficiently small g_f , I can show that the second effect dominates. This follows from an intuitive continuity argument, since on shutting down familiarity growth the model reverts to a standard one in which the first effect is zero.

Proposition 4. *There exists $\bar{g}_f > 0$ such that for all $g_f \in [0, \bar{g}_f]$, the entry value is monotone decreasing in τ over the domain $(0, \eta/\bar{v})$ over which BGP with positive entry exists. Hence, if a BGP with positive entry exists, it is unique.*

When I calibrate the model, I numerically determine the upper bound $\bar{g}_f$ that ensures uniqueness. In all calibration exercises uniqueness holds by a wide margin.

4 How does ageing affect growth?

This section studies how ageing, modelled as a decrease in the death hazard μ , affects growth through adoption frictions for new products and firm entry. It starts by showing that, within the model developed so far, ageing reduces the value of opening a new firm, which implies that ageing reduces entry of new firms and products and slows down growth. The section then takes the model to the data, describing the calibration strategy and deriving the main quantitative results of the paper, showing that ageing explains 32.1% of the growth slowdown between the 1980s and the 2010s.

4.1 Ageing reduces growth

I have already shown that for a product of age zero, the familiarity index is increasing in the death hazard, $\partial_\mu \mathcal{F}(0) > 0$: in an older society (with lower μ), all else equal, a newly introduced product sells less since consumers are more familiar with existing products. This follows

from the fact that consumers who have lived longer have had more time to become familiar with existing products. However, it is not only sales upon entry that matter for potential entrants: they also expect sales to be higher in the future (so that the entry condition depends on $\kappa(0)\mathcal{F}(0) = \widehat{\mathcal{F}}(\rho + (1 + g_q)\tau; \mu, \tau)$, not on $\mathcal{F}(0)$ alone). To show that ageing lowers entry and growth, I use the implicit function theorem. Since the entry condition is

$$\widehat{\mathcal{F}}(\rho + (1 + g_q)\tau; \mu, \tau) = \frac{\bar{v}}{1 + g_q},$$

the derivative of the creative destruction rate with respect to the death hazard is

$$\frac{d\tau}{d\mu} = -\frac{\partial_\mu \widehat{\mathcal{F}}}{\frac{d}{d\tau} \widehat{\mathcal{F}}} \quad (38)$$

whenever the denominator is different from zero. I established above that for $g_f \leq \bar{g}_f$ the denominator is negative and different from zero in the relevant range in which the equilibrium exists. Next, I show that the direct effect of ageing is to lower the entry value, obtaining the growth result.

Proposition 5. *Let $g_f < \min\{(1 + g_q)\tau, \bar{g}_f\}$. Then ageing lowers the entry value, $\partial_\mu \widehat{\mathcal{F}} > 0$. This implies that ageing lowers entry and the creative destruction rate.*

Proof. In the proof I hold τ fixed throughout, and show that the entry value increases in the death hazard μ for each level of creative destruction. The key observation on which the proof builds is that $e^{-(\rho + (1 + g_q)\tau)u}$ factorises into an impatience term $e^{-\rho u}$ and the quality-tilted discount rate $e^{-(1 + g_q)\tau u}$. Then I can write the entry value as

$$\widehat{\mathcal{F}}(\rho + (1 + g_q)\tau; \mu, \tau) = \int_0^\infty e^{-\rho u} \mathcal{F}(u) e^{-(1 + g_q)\tau u} du, \quad (39)$$

where the quality-tilted discount $e^{-(1 + g_q)\tau u}$ is, up to the factor $(1 + g_q)\tau$, the quality-weighted density of product ages against which $\mathcal{F}$ averages to one (23). Exploiting the representation in the proof of Proposition 1(iii), I can show that the last integral is dominated by an integrable function $C e^{-(\rho + (1 + g_q)\tau - g_f)u}$, so that I may differentiate (39) under the integral sign, obtaining

$$\partial_\mu \widehat{\mathcal{F}} = \int_0^\infty e^{-\rho u} \partial_\mu \mathcal{F}(u) e^{-(1 + g_q)\tau u} du. \quad (40)$$

What is crucial here is that only the familiarity index $\mathcal{F}$ and not the weight $e^{-(1 + g_q)\tau u}$ depends on μ . Two properties of $\partial_\mu \mathcal{F}$ pin down the sign of this integral. First, ageing is *purely redistributive* across product ages: the weight does not depend on μ while the normalisation (23) holds

identically in μ , so differentiating it in μ yields the conservation identity

$$\int_0^\infty \partial_\mu \mathcal{F}(u) e^{-(1+g_q)\tau u} du = 0, \quad (41)$$

the interchange being justified by the same domination. Second, by Proposition 1(iii), $\partial_\mu \mathcal{F}$ is strictly positive on $[0, a^{p*})$ and strictly negative on (a^{p*}, ∞) . Subtracting from (40) the identity (41) multiplied by $e^{-\rho a^{p*}}$,

$$\partial_\mu \widehat{\mathcal{F}} = \int_0^\infty (e^{-\rho u} - e^{-\rho a^{p*}}) \partial_\mu \mathcal{F}(u) e^{-(1+g_q)\tau u} du. \quad (42)$$

Since $\rho > 0$, the tilt factor $e^{-\rho u} - e^{-\rho a^{p*}}$ changes sign in the same place a^{p*} and in the same direction as $\partial_\mu \mathcal{F}$: both are strictly positive for $u < a^{p*}$ and strictly negative for $u > a^{p*}$. The integrand of (42) is therefore strictly positive for every $u \neq a^{p*}$, which implies $\partial_\mu \widehat{\mathcal{F}} > 0$.

Finally, since $g_f \leq \bar{g}_f$, Proposition 4 implies that the denominator of (38) is strictly negative on the domain over which a BGP with positive entry exists, so $d\tau/d\mu > 0$: ageing, a fall in μ , reduces entry and the creative destruction rate. $\square$

The economics of the result is as follows: ageing only reassigns demand along the quality-weighted distribution, from young to old products. The quality-weighted distribution discounts old products for two reasons: one is that they are more likely to have been replaced (rate τ) and the other is that they sell less anyway because their quality falls behind the average at rate $g_q\tau$. However, a prospective entrant values young products relatively more: in addition to the probability of being destroyed τ , the opportunity cost r exceeds aggregate growth by $\rho > 0$, since households are impatient. Then, ageing lowers the entry value.³⁷

As a result, the total effect of ageing on the sales of a new product is stronger than the direct effect through mortality alone, since a lower entry rate implies existing products have more time to acquire familiarity.

Under the hypotheses of the proposition above, the total effect of ageing on new-product demand is

$$\frac{d\mathcal{F}(0)}{d\mu} = \underbrace{\partial_\mu \mathcal{F}(0)}_{>0 \text{ (Prop. 1(iii))}} + \underbrace{\partial_\tau \mathcal{F}(0)}_{>0 \text{ (Prop. 1(iv))}} \cdot \underbrace{\frac{d\tau}{d\mu}}_{>0} > \partial_\mu \mathcal{F}(0) > 0.$$

A change in entry has a linear effect on growth, and through $r = \rho + \phi + g$ on the interest rate, so that it is immediate to show that ageing reduces growth and the interest rate

$$\frac{dg}{d\mu} = \frac{g_q}{\eta - 1} \frac{d\tau}{d\mu} > 0, \quad \frac{dr}{d\mu} = \frac{dg}{d\mu} > 0.$$

³⁷If $\rho = 0$, ageing is indeed neutral on the entry value.

The age profile of expenditure. The results of this section are written for expenditure that does not change with age, but they hold for any expenditure profile, with the spending hazard $\tilde{\mu} = \mu - \phi$ in place of μ and the discount rate $\tilde{\rho} = \rho + \phi$ in place of ρ : the profile enters the free-entry condition (33) only through these two rates. The economy with profile ϕ at (μ, ρ) therefore has the same rates of creative destruction, growth and interest as the economy with constant expenditure at $(\tilde{\mu}, \tilde{\rho})$: the profile moves ϕ from the death hazard to the discount rate, leaving the sum $\mu + \rho$ unchanged. Ageing, a fall in μ at a given ϕ , lowers $\tilde{\mu}$ one for one and leaves $\tilde{\rho}$ unchanged, so it lowers entry and growth as shown above. At a given entry rate, a steeper profile has two effects on the value of entering,

$$\frac{d}{d\phi} \widehat{\mathcal{F}}(\tilde{\rho} + (1 + g_q)\tau; \tilde{\mu}, \tau) = \underbrace{\frac{\partial \widehat{\mathcal{F}}}{\partial \sigma}}_{\text{discounting}} - \underbrace{\frac{\partial \widehat{\mathcal{F}}}{\partial \tilde{\mu}}}_{\text{older spending}}.$$

Through discounting, it raises the gap between the return and growth, $r - g = \rho + \phi$, so the entrant weighs its early, unfamiliar years more heavily against the familiar maturity that follows. Through demand, it acts like population ageing: it shifts spending towards older households, who are more familiar with existing products. As ϕ falls towards $-\rho$, the entry value tends to $1/[(1 + g_q)\tau]$ whatever the demographics, by the normalisation (23): the ageing channel switches off. Moving expenditure towards older households therefore lowers entry and growth much as population ageing does. On the demand side, what matters for innovators is the age of expenditure, not of the population: the mean age of the dollars spent, $1/\tilde{\mu}$, rather than the mean age of households, $1/\mu$.

4.2 Calibrating the model

The quantitative experiment proceeds as follows: I calibrate the model to the period 1980–1989, the “young” era, and then increase the length of life and change the age profile of expenditure to match data from the period 2010–2019, the “old” era, keeping all other parameters fixed. I then ask how much of the observed decline in growth the model explains. I use decade averages to reduce the influence of annual fluctuations, and end the comparison in 2019 to exclude the pandemic.

The model is extremely parsimonious: compared to the standard quality ladder model, it only features three additional parameters, the death hazard μ , the slope of the expenditure profile ϕ and the familiarity growth rate g_f (the baseline level of familiarity f_0 doesn’t affect allocations, only the level of utility, and can be freely normalised). Overall, I have a set of seven parameters $\{\rho, \mu, \phi, \eta, g_q, g_f, \bar{v}\}$ to calibrate. I set $\eta = 3$, within the range of the elasticities of substitution estimated by Broda and Weinstein (2006, 2010), and $\rho = 0.05$, which is a standard value

and implies a return of 7% at 2% yearly growth, broadly consistent with the long-run return on equities. Given η and the young-era values for entry τ and growth g , I recover the quality step from (20) as $g_q = (\eta - 1)g/\tau = 0.351$.

I set μ in both periods using the mean adult age in the Census, measured as years above 18, which was 25.4 years in the young era and 29.1 in the old era. I use adult age since in the model what matters is how long agents live as independent economic actors. Since mean age in the model is $1/\mu$, I take the inverse of mean adult age in each year and average it within each era, obtaining $\mu^{\text{young}} = 0.0394$ and $\mu^{\text{old}} = 0.0344$. Table 4 reports the macroeconomic and demographic moments.

Table 4: Calibration moments

Moment	Period	Value	Source
GDP p.c. growth	1980–1989	2.10%	FRED
Firm entry rate	1980–1989	11.99%	Census BDS
Mean adult age (18+)	1980–1989	25.4 yr	Census
Mean adult age (18+)	2010–2019	29.1 yr	Census

Calibrating the expenditure profile. I showed in Section 4.1 that what matters for innovators is the age of the dollars spent, not of households. In the model, the mean adult age of households is $1/\mu$, while that of expenditure, weighting each household by its spending, is $1/(\mu - \phi)$. I measure both in the Consumer Expenditure Survey (CEX) and recover the slope of the expenditure profile from the difference between their inverses,

$$\phi = \frac{1}{\bar{a}} - \frac{1}{\tilde{a}},$$

where $\bar{a}$ is the mean adult age and $\tilde{a}$ the mean adult age with each adult weighted by consumption, both in years above 18. Since households in the model are individuals, I split the consumption of each household in the CEX equally among its members aged 18 and over.³⁸ In the 1980s, $\bar{a} = 25.5$ and $\tilde{a} = 24.4$ years: spending per adult fell slightly with age, and $\phi^{\text{young}} = -0.0019$. In the 2010s, the ranking had reversed, with $\bar{a} = 29.3$ and $\tilde{a} = 29.7$ years, so that $\phi^{\text{old}} = 0.0005$. I keep ρ fixed across eras, so that the gap between the return and growth, $r - g = \rho + \phi$, rises from 4.81% to 5.05%.

³⁸Consumption is total expenditure minus cash contributions and payments for personal insurance and pensions. I compute ϕ in each year and average it within each era. The 1980s average uses the surveys of 1980–1981 and 1984–1989.

Calibrating familiarity growth. I calibrate g_f using the age profile of spending on new products in the HMS. The model features an intensive margin of adoption: all households consume all products, but older households spend less on new ones because they are more familiar with existing goods. The spending share therefore provides the empirical counterpart to the mechanism in the model.

Let $S(a^h)$ be the share of expenditure that a household of adult age a^h spends on products in their first year. For households older than these products, familiarity with each new product is the same across households. Using individual demand (8), the only term in the spending share that varies with household age is then the ideal price index. The ratio of spending shares between a household of age a^h and one of reference age a_0^h is

$$\frac{S(a^h)}{S(a_0^h)} = \left(\frac{P(a^h)}{P(a_0^h)} \right)^{\eta-1} = \frac{(1 + g_q)\tau - g_f e^{-[(1+g_q)\tau - g_f]a_0^h}}{(1 + g_q)\tau - g_f e^{-[(1+g_q)\tau - g_f]a^h}}. \quad (43)$$

The ratio does not depend on the expenditure profile, which changes how much a household spends but not how it divides its spending across products. Given the young-era entry rate and the calibrated quality step, $(1 + g_q)\tau$ is fixed, so that g_f alone determines the relative age profile. A larger g_f makes the profile steeper: older households have accumulated more familiarity with existing products and devote a smaller share of spending to new ones.


I use the age effects of regression (1), estimated with all its terms on the years 2005–2024 and plotted in Figure 2. New products are those introduced in the current year, and expenditure is total household spending on the packaged goods covered by the panel. The profile consists of 62 points: one for each age between 26 and 85, and two for the pooled groups aged 25 and below and 86 and above, with age 30 as the reference. For each point b , I express the estimated difference $\hat{\beta}_b$ relative to the average spending share of households aged 30, $\bar{S}_0$, obtaining $\hat{s}_b = 1 + \hat{\beta}_b / \bar{S}_0$.

I evaluate the model at each age, and at the projection-weighted mean age of each pooled group, measured in completed years above 18, and choose g_f to minimise the sum of squared differences between the model and empirical profiles,

$$\hat{g}_f = \arg \min_{g_f > 0} \sum_b \left[\frac{S(a_b^h; g_f)}{S(a_0^h; g_f)} - \hat{s}_b \right]^2.$$

Each point receives equal weight; the reference age contributes zero since both profiles equal one there. The fitted value is $g_f = 0.099$. Although I do not impose it, the estimate satisfies $g_f < (1 + g_q)\tau = 0.162$, as Proposition 5 assumes. Figure 5 compares the resulting model profile with the data.

Figure 5: The age profile of spending on new products and the calibrated model



In the calibrated case $g_f < (1 + g_q)\tau$, the profile (43) converges as households age, since their price index does (Section 3.7):

$$\lim_{a^h \rightarrow \infty} \frac{S(a^h)}{S(a_0^h)} = 1 - \frac{g_f}{(1 + g_q)\tau} e^{-[(1 + g_q)\tau - g_f]a_0^h} = 0.706 .$$

Very old households thus spend 29.4% less on new products than households aged 30. The model is already close to this limit at 60, when the share is 26.2% below its level at 30, and at 85 it is 28.8% below it: after 60, it falls by only about three more points. Over the same ages, the data fall from about 26% to 43% below the level at 30, which is why the model profile in Figure 5 is flatter than the data among older households. The convergence also means that perpetual youth is not an issue for the calibration. With a constant death hazard, some households live counterfactually long lives. Since the profile converges, however, these households spend on new products almost exactly as households aged 85 do, so that the unrealistic tail of the age distribution matters little for the demand for new products.

Given g_f , I obtain the research cost scale $\bar{v}$ from the free-entry condition (33), evaluated at the young-era entry rate. This gives $\bar{v} = 5.58$ and completes the calibration. Table 5 reports the parameters. In the ageing experiment, I then replace μ^{young} and ϕ^{young} with μ^{old} and ϕ^{old} and solve the free-entry condition again, keeping all other parameters fixed. The old-era entry and growth rates are outcomes of the experiment.

Table 5: Calibrated parameters

Parameter	Value	Mapping
μ^{young}	0.0394	Mean adult age (18+), 1980–1989
μ^{old}	0.0344	Mean adult age (18+), 2010–2019
ϕ^{young}	−0.0019	Age profile of spending per adult (CEX), 1980–1989
ϕ^{old}	0.0005	Age profile of spending per adult (CEX), 2010–2019
η	3	Broda and Weinstein (2006, 2010)
ρ	0.05	Standard in the literature, implies $r = 7\%$ at $g = 2\%$
g_q	0.351	Growth rate, 1980–1989
g_f	0.099	Age profile of spending on new goods
$\bar{v}$	5.58	Entry rate, 1980–1989

4.3 How much of the growth slowdown does ageing explain?

Table 6 reports the result of the ageing experiment. The young era is targeted, so the model reproduces the growth rate of 2.10% in 1980–1989 by construction. With the death hazard and the expenditure profile of the old era, growth falls to 1.95%, against 1.64% in the data: ageing explains 32.1% of the growth slowdown.

Two forces are at work: households live longer, and spending moved towards older households. To isolate the first, I recalibrate the model with a flat expenditure profile, $\phi = 0$, in both eras, so that only the death hazard changes. Growth then falls to 2.05%, and ageing explains 11.1% of the slowdown, about a third of the total effect.

The ageing of expenditure lowers growth through the two channels I described in Section 4.1. Spending aged faster than the population: between the two eras, the mean adult age of expenditure, $1/(\mu - \phi)$, rose by 5.3 years, against 3.7 years for households. A steeper expenditure profile also raises the gap between the return and growth, $r - g = \rho + \phi$, from 4.81% to 5.05%, so entrants discount more heavily the years in which their products are still unfamiliar.

Ageing is therefore a sizeable force behind the growth slowdown, and two thirds of its effect come from the interaction of ageing with expenditure shifting to older consumers.

5 Along the transition

Ageing is a slow process that takes decades to happen, so transition between steady states should be important. This section studies the transition between the young and the old balanced growth paths.³⁹ Although the entry technology is linear, the economy does not jump between balanced growth paths: even after a fall in mortality, the age distribution and the fa-

³⁹Appendix D describes the algorithm used to compute the transition.

Table 6: Share of the growth decline explained by ageing

	Data	Model	
		$\phi = 0$	Total
<i>GDP per capita growth g</i>			
Young (1980–1989)	2.10%	2.10%	2.10%
Old (2010–2019)	1.64%	2.05%	1.95%
Share of the decline explained	—	11.1%	32.1%

Notes. Model, $\phi = 0$: flat expenditure profile in both eras, so that only the death hazard changes. Total: the death hazard and the expenditure profile both change. Each version of the model is calibrated to the young era.

miliarity stock it carries are slow-moving state variables, and entry responds to them as they evolve.⁴⁰

In this exercise, I use the demographic path the United States actually experienced, together with the observed ageing of expenditure. I take the census mean adult age year by year from 1980, and the expenditure profile ϕ_t from the CEX, computed as in Section 4.2 for every year and smoothed with a centred 11-year moving average; after 2024, ϕ stays at its 2024 value. Since demand behaves as if households died at the spending hazard $\tilde{\mu} = \mu - \phi$, I feed the model the mean adult age of expenditure, and recover the sequence of spending hazards $\{\tilde{\mu}_t\}$ that reproduces it.⁴¹ Entrants discount profits at $r_t - g_t = \rho + \phi_t$, which rises from 4.80% in 1980 to 5.02% from 2024. Figure 6 shows the two inputs. Expenditure aged faster than the population: between 1980 and 2025, the mean adult age of expenditure rose from 41.9 to 48.7 years, against 43.1 to 48.5 for people. I calibrate the initial condition to 1979, the year before the simulation starts: the economy starts on the balanced growth path associated with the 1980 age distribution of expenditure and the 1979 rates of firm entry, 12.82%, and GDP per capita growth, 2.02%. These rates give $g_q = 0.314$; refitting the age profile of Section 4.2 at them gives $g_f = 0.104$, and free entry gives $\bar{v} = 5.27$. The other parameters are those of Table 5. I then solve for the perfect-


⁴⁰The path converges to the new balanced growth path, for changes in mortality in a neighbourhood of the initial BGP. The value of entering today depends on entry at every other date: past entry raises it, since it leaves younger products on the market, with which consumers are less familiar; future entry lowers it, since it shortens the expected life of the new product. Without familiarity growth, past entry plays no role and entry jumps after a permanent change in parameters, as in standard quality-ladder models. With familiarity, past entry is a state variable. The linearised entry condition has exactly one convergent solution, regardless of the inherited initial state, if and only if the Fourier transform of these effects never vanishes and has a winding number equal to zero (Krein, 1962). This condition extends Blanchard and Kahn (1980) to infinitely many leads and lags; a zero winding number intuitively implies that there must be as many stable roots as predetermined variables. Auclert et al. (2025) apply this analysis to sequence-space Jacobians. I check that the condition holds at the calibration. The largest stable root plays the role of the stable eigenvalue in the neoclassical growth model: it determines the speed of convergence in the long run, and thus the half-life of the transition.

⁴¹Mean age and the hazard are linked by the renewal identity $\dot{m}(t) = 1 - \tilde{\mu}_t m(t)$ for the mean adult age of expenditure m , which I invert to obtain $\tilde{\mu}_t$ from the path of m .

foresight equilibrium path of entry from 1980 onwards.

Beyond 2025 the demographic path is no longer observed, and I consider two scenarios. In the first, ageing stops: the mean adult age of expenditure is held at its 2025 value. In the second, ageing continues as projected by the Census Bureau: the population follows the 2023 middle-immigration projections until 2100, with ϕ at its 2024 value, and the mean adult age of expenditure is held constant thereafter.

Figure 6: Ageing of the population and of expenditure




Left: ϕ_t in the CEX, in percent per year of age, year by year and as a centred 11-year moving average, the series used in the model. Right: the mean adult age of people (Census) and of dollars spent, implied by the census and the smoothed ϕ_t ; the model is fed the second.

Figure 7 reports the result. The left panel shows the demographic input, the mean adult age of people and of dollars spent. The middle panel shows the growth rate, and the right panel GDP per capita relative to the 1979 balanced growth path, the path the economy would have followed had the age structure of its expenditure stayed at that of 1980.

Two main insights emerge from the transition analysis.

First, ageing initially causes an expansion: growth rises from 2.02 to 2.05 percent per year in the early 1980s, and remains above its initial value until 1990. The reason is that while mortality falls immediately, the age distribution and the familiarity stock are slow-moving, so there is a premium to entering early. An innovator entering at the outset sells at introduction exactly as much as on the initial balanced growth path but sells comparatively more in the future: by the time the population has actually aged, its product will be one of the older, familiar lines whose demand ageing raises (Proposition 1(iii)). At the same time, since innovation will slow

Figure 7: The transition along the observed demographic path



Solid: the census path and the Census Bureau projections to 2100, then frozen. Dashed: the mean adult age of expenditure frozen at its 2025 value. Left: the mean adult age of people (Census) and of dollars spent. Middle: the growth rate. Right: GDP per capita relative to the 1979 balanced growth path, the economy that does not age, in percent. Grey dashed lines: the 1979 balanced growth path.

along the transition, a product introduced today is displaced later in expectation, which raises the present value of innovating today. Innovation booms until the age distribution itself turns over. Two related results also follow. The expansion is increasing in the size of the demographic change to come, so a surge in growth can be a symptom of worse demographic news ahead. And in those years growth looks healthier than in the pre-ageing economy, even though the fundamentals of innovation have already deteriorated. Hence, in the model growth does not merely lag the demographic change; it initially points in the wrong way. The expansion comes from anticipating that consumers will become more familiar with existing products: in the appendix, when households become older at once but keep their familiarity, the same expansion appears as long as the gap between the return and growth stays at its young-era value (Figure 17), and it almost disappears when innovators do not foresee future ageing and update their beliefs every year (Figure 18).

Second, the transition is slow and the decline is back-loaded. Growth falls back below its 1979 level only in 1990, and reaches 1.84 percent in 2025. If ageing stops in 2025, growth settles at 1.85 percent. If ageing continues as projected, growth in 2025 is higher, at 1.87 percent, since innovators anticipate the ageing to come, as in the 1980s; it then keeps falling, to 1.74 percent by 2100, so that by 2025 the economy has gone through only about half of its eventual decline. Output tells the same story. Had the age structure of expenditure stayed at that of 1980, GDP

per capita would be 3.5% higher in 2025 (3.1% if ageing continues). Even if ageing stops, the gap reaches 8.0% by 2050 and 13.5% by 2080, since the two economies grow at permanently different rates; with the projections, it reaches 15.6% by 2080. The model thus offers a mechanism for why a demographic turnaround that began around 1980 shows up in growth only from the 1990s onwards, and implies that most of the output cost of ageing is still to come.

6 How does ageing affect research subsidies?

This section studies how ageing changes the effectiveness of research subsidies, a policy tool that arises naturally in models of innovation. Under a research subsidy, governments try to stimulate growth by paying a part of the research costs. Formally, entrants pay the fraction $1 - s$ of the entry cost and the government the remaining fraction s , financed by a lump-sum tax on households proportional to their income $\phi(a^h)$; $s < 0$ is an entry tax. Since the tax follows the age profile of income, the age profile of expenditure does not change, the Euler equation and the static allocation are untouched, and the resource cost of research is still $\bar{v}\tau Y/\eta$: the subsidy is a transfer, and the only equilibrium object it moves is free entry, which becomes

$$\widehat{\mathcal{F}}(\bar{\rho} + (1 + g_q)\tau; \bar{\mu}, \tau) = (1 - s) \frac{\bar{v}}{1 + g_q}. \quad (44)$$

When $g_f \leq \bar{g}_f$, the entry value is strictly decreasing in τ (Proposition 4). Differentiating (44) then gives the effect of the subsidy on entry at the subsidised equilibrium and, since $g = g_q\tau/(\eta - 1)$, on growth:

$$\frac{d\tau}{ds} = \frac{\bar{v}}{1 + g_q} \left| \frac{d}{d\tau} \widehat{\mathcal{F}} \right|^{-1} > 0, \quad \frac{dg}{ds} = \frac{g_q}{\eta - 1} \frac{d\tau}{ds}. \quad (45)$$

The subsidy is powerful when the value of entry falls slowly as entry expands, since by (45) a flat entry value lets a small subsidy move entry a lot. By (37), the slope of the entry value is the net of the two forces of Section 3: more entry shortens the expected life of each product, but it also resets the familiarity of consumers with existing products, which helps the products that enter next.

How does ageing change the effectiveness of research subsidies? For this section write $\widehat{\mathcal{F}}(\tau, \bar{\mu})$ for the composed entry value $\widehat{\mathcal{F}}(\bar{\rho} + (1 + g_q)\tau; \bar{\mu}, \tau)$, with subscripts for its derivatives. Differentiating along the ageing experiment splits the answer into two channels.

Proposition 6. *Along the equilibrium path $\tau(\tilde{\mu})$, at any implementable s ,*

$$\frac{d}{d\tilde{\mu}} \frac{d\tau}{ds} = \frac{\bar{v}}{(1 + g_q) \widehat{\mathcal{F}}_\tau^2} \left[\underbrace{\widehat{\mathcal{F}}_{\tau\tilde{\mu}}}_{\text{familiarity}} + \underbrace{\widehat{\mathcal{F}}_{\tau\tau} \frac{d\tau}{d\tilde{\mu}}}_{\text{relocation}} \right], \quad \frac{d\tau}{d\tilde{\mu}} = -\frac{\widehat{\mathcal{F}}_{\tilde{\mu}}}{\widehat{\mathcal{F}}_\tau}. \quad (46)$$

The decomposition follows from differentiating (45) in $\tilde{\mu}$ along the equilibrium path, with $d\tau/d\tilde{\mu}$ obtained by implicit differentiation of (44).

The familiarity term is the effect of ageing on the slope of the entry value at fixed destruction, and it runs through the familiarity profile of the households who spend: at fixed τ and $\tilde{\rho}$, the spending hazard only enters the entry value through this. The relocation term is the movement of the equilibrium along the curve to a point of different slope, the entry decline of Proposition 5 interacting with the curvature of the entry value.

Table 7 reports the effectiveness at the calibration of Section 4, at each decade's decentralised equilibrium. Ageing makes the instrument more powerful: a marginal point of subsidy buys 0.22 points of entry in the young era and 0.26 in the old (+16.2 percent); in growth units, a subsidy covering ten percent of entry costs buys 0.42 percentage points of growth in the young economy and 0.47 in the old. Integrated along the path between the two eras, the familiarity channel of (46) accounts for 86% of the increase. Relocation works the other way and subtracts 19%: entry is lower in the old era, where the entry value falls faster as entry expands. The remaining third comes from discounting. The shift of expenditure towards older households also raises $\tilde{\rho} = r - g$, from 4.81% to 5.05%, which flattens the entry value at a given entry rate and adds to (46) a term analogous to the familiarity one.

Table 7: Ageing and the effectiveness of the entry subsidy

	1980s	2010s	Change
Marginal effectiveness $d\tau/ds$ at $s = 0$	0.223	0.259	+16.2%
Growth bought by $s = 0.1$	+0.42 pp	+0.47 pp	+11.0%

Why does ageing make the subsidy more powerful? Familiarity is reset in two ways. Households die and are replaced by newborns who begin their life unfamiliar with products, and products are replaced by new ones that no household knows. A research subsidy only speeds up the second reset channel. In a very young society, death does most of the work: households do not live long enough to become very familiar with any product, whatever the pace of creative destruction. Faster creative destruction then barely changes the familiarity that entrants compete against, and the subsidy mainly shortens the life of each product. In an old society, households outlive the products they buy, and familiarity is curbed mainly by creative destruction.

Each subsidised entrant then lowers the familiarity of consumers with existing products, and the products that enter next sell more. Ageing moves the economy towards this second case. At the calibration, the reset of familiarity, the first term of (37), offsets 52% of the effect of shorter product lives, the second term, in the 1980s, and 62% in the 2010s. On net, the value of entry falls more slowly as entry rises, and the same subsidy buys more entry.

7 Does ageing improve utility and welfare?

The previous sections showed that ageing slows down entry and growth. It is then natural to ask what effect ageing has on welfare. To answer this question, though, we must take a stance on an issue that the positive analysis could leave open: does familiarity really raise utility? Until now, this did not matter. Familiarity enters the consumption aggregator (5), but in Sections 4 to 6 its only role is to tilt demand towards the products a household knows well. Habits, a cost of searching among products or a lack of information about new ones would generate the same demand structure, and positive results. Familiarity could then be read as a modelling device. For welfare instead, if familiarity is part of preferences, a household that knows the products it buys enjoys the same basket more, and an older population, which knows more products, is better off. If familiarity is a friction, the tilt towards familiar products can be mistake, and the familiarity built up by ageing is a cost. In the remainder of this section I take the first view: consumers enjoy the products they know, from a brand they trust to software they have learnt to use. In Section 7.3, I discuss how the quantitative results change if familiarity doesn't increase utility.

I obtain three main results. First, at any given rate of creative destruction, an older population is more familiar with the products it buys, and this raises welfare. Second, for an individual the gain is larger, since she also values living longer. Third, along the transition welfare depends on a race between the gains in familiarity, which raise the level of utility, and the growth loss, which compounds forever.

7.1 Welfare

I define social welfare as the discounted flow of the utilities of the households alive at each date,

$$\mathcal{W} = \int_0^\infty e^{-\rho t} \int_0^N \log C_{it} \, di \, dt . \quad (47)$$

Welfare depends on how expenditure and familiarity change with age. A household of age

a^h spends $\phi(a^h)E$ by (6), and by (11) it consumes

$$C(a^h) = \phi(a^h) E \mathcal{F}^h(a^h)^{\frac{1}{\eta-1}}, \quad (48)$$

where $\mathcal{F}^h$ is the average familiarity (24): it starts at the reset level $\mathcal{F}^h(0) = f_0$ and rises with age. Two households with the same budget buy the same quantity of goods, but allocate it differently across products. The older one is more familiar with existing products, so it obtains more consumption services. Averaging the log of (48) over the age distribution of the population (2), per-capita flow utility at any date is

$$\frac{1}{N} \int_0^N \log C_{it} \, di = \log E_t + \mathbb{E}[\log \phi(a^h)] + \frac{1}{\eta-1} \mathbb{E}[\log \mathcal{F}_t^h(a^h)], \quad (49)$$

The three terms are average expenditure, the dispersion of expenditure across ages and the average log familiarity of the living population. The income profile averages to one across the population, so by the concavity of the log the second term is negative unless $\phi = 0$. On a balanced growth path it equals $\log(1 - \phi/\mu) + \phi/\mu$.

Since familiarity multiplies consumption, one unit of consumption is worth different amounts to different households. I therefore express welfare comparisons in a metric that is common to everyone: a uniform share of consumption. For two scenarios A and B , the consumption equivalent of moving from B to A is the share x by which every household's consumption of every good must rise, at every date of scenario B , to make welfare equal in the two scenarios. The aggregator (5) is homogeneous of degree one, so this change raises every household's flow utility by the same $\log(1+x)$, whatever its age or familiarity, and

$$\text{CE} = \exp \left\{ \frac{\rho}{N} (\mathcal{W}^A - \mathcal{W}^B) \right\} - 1. \quad (50)$$

The metric values familiarity in units of consumption. It is also free of normalisations: the baseline familiarity f_0 , the initial technology and the units of goods all cancel in $\mathcal{W}^A - \mathcal{W}^B$. Section 7.2 introduces the counterpart for an individual newborn.

I first compare balanced growth paths that start from the same technology $\bar{Q}_0$. On a balanced growth path, per-capita expenditure is $E_t = (1 - \bar{\nu}\tau/\eta) \bar{Q}_t$ by (34), technology grows at the rate $g = g_q\tau/(\eta-1)$, and the familiarity profile (24) does not change over time. Substituting (49)

into (47), for a given expenditure profile ϕ ,

$$\begin{aligned} \frac{\rho \mathcal{W}}{N} = & \log \bar{Q}_0 + \underbrace{\frac{\log f_0}{\eta - 1}}_{\text{resources}} + \underbrace{\log \left(1 - \frac{\bar{v} \tau}{\eta}\right) + \log \left(1 - \frac{\phi}{\mu}\right) + \frac{\phi}{\mu}}_{\text{life cycle}} \\ & + \underbrace{\frac{1}{\eta - 1} \mathbb{E} \left[\log \frac{\mathcal{F}^h(a^h)}{f_0} \right]}_{\text{familiarity}} + \underbrace{\frac{\tau g_q}{(\eta - 1) \rho}}_{\text{growth}}. \end{aligned} \quad (51)$$

Besides constants, welfare has four terms. The resource term is the share of output that is consumed rather than spent on entry. The life-cycle term is the cost of spreading expenditure unevenly across ages. The familiarity term is the average log familiarity of the living population. The growth term, g/ρ , is the discounted value of growth. Creative destruction enters three of the terms: it costs resources, erodes familiarity and generates growth. Given τ , the death hazard enters only the familiarity and life-cycle terms.

The familiarity term has a closed form, tied to the familiarity index of a new product.

Lemma 3. *On a balanced growth path,*

$$\mathbb{E} \left[\log \frac{\mathcal{F}^h(a^h)}{f_0} \right] = \frac{g_f - (1 + g_q) \tau (1 - \mathcal{F}_\mu(0))}{\mu}, \quad (52)$$

where $\mathcal{F}_\mu(0) \equiv \int_0^\infty \mu e^{-\mu a^h} f_0 / \mathcal{F}^h(a^h) da^h$ is the familiarity index of a new product with every household weighted equally: the index $\mathcal{F}(0)$ of Section 3.7 computed at the death hazard μ instead of the spending hazard $\tilde{\mu}$.

The lemma requires no restriction on the parameters: unlike the tail of the index (Proposition 1), the log average is always finite, since $\log(\mathcal{F}^h(a^h)/f_0) \leq g_f a^h$. Substituting the lemma into (51) gives welfare on a balanced growth path in closed form:

$$\begin{aligned} \frac{\rho \mathcal{W}}{N} = & \log \bar{Q}_0 + \underbrace{\frac{\log f_0}{\eta - 1}}_{\text{resources}} + \underbrace{\log \left(1 - \frac{\bar{v} \tau}{\eta}\right) + \log \left(1 - \frac{\phi}{\mu}\right) + \frac{\phi}{\mu}}_{\text{life cycle}} \\ & + \underbrace{\frac{1}{\eta - 1} \frac{g_f - (1 + g_q) \tau (1 - \mathcal{F}_\mu(0))}{\mu}}_{\text{familiarity}} + \underbrace{\frac{\tau g_q}{(\eta - 1) \rho}}_{\text{growth}}. \end{aligned} \quad (53)$$

Equation (53) also shows a cost of innovation that innovators do not internalise. In the free-entry condition (33), an entrant values its own familiarity path $\widehat{\mathcal{F}}$. The familiarity that its arrival destroys, $(1 + g_q)(1 - \mathcal{F}_\mu(0))$ per innovation, enters no firm's value. Section 8 studies this externality.

In this comparison, longer lives have no direct effect on welfare. Births adjust one for one, so (47) counts the same N households at every date, whatever μ . A fall in mortality turns many short lives into fewer long ones, and welfare is neutral to this change, apart from the life-cycle term.

Fix τ . At a lower μ , the age distribution (2) is first-order stochastically larger. Average familiarity rises with age, and so does its log. The familiarity term of (51) therefore rises strictly as μ falls, for every $g_f > 0$. The life-cycle term moves the other way: at a given ϕ , its derivative in μ is $\phi^2 / [\mu^2(\mu - \phi)] \geq 0$, because longer lives spread a given slope of expenditure over more ages. This effect is of second order in ϕ and vanishes at $\phi = 0$.

Proposition 7. *For every fixed $\tau \in [0, \eta/\bar{v})$, the familiarity term of (51) is strictly decreasing in μ . When $\phi = 0$, so is welfare $\mathcal{W}$, and its maximum over the rate of creative destruction, $\mathcal{W}^*(\mu) \equiv \max_{\tau \in [0, \eta/\bar{v})} \mathcal{W}$, is strictly decreasing in μ : across balanced growth paths, an older population achieves strictly higher welfare under optimal steady-state policy, and under any policy that pins the rate of creative destruction as μ falls.*

Across the two eras, however, τ is not fixed: ageing lowers entry, $d\tau/d\mu > 0$. Along the equilibrium, ageing therefore has two effects on welfare:

$$\frac{d\mathcal{W}}{d\mu} = \underbrace{\frac{\partial \mathcal{W}}{\partial \mu}}_{\substack{<0 \text{ at } \phi=0 \\ (\text{Prop. 7})}} + \underbrace{\frac{\partial \mathcal{W}}{\partial \tau} \cdot \frac{d\tau}{d\mu}}_{>0}. \quad (54)$$

The first term is the familiarity gain from ageing, net of the life-cycle term. The second is the effect of slower entry, and its sign is ambiguous. If equilibrium entry is at or above the rate that maximises welfare, slower entry is itself welcome. If entry is too low, slower entry deepens the shortfall. Which effect dominates is a quantitative question, which I answer in Section 7.3.

7.2 Individual utility and longer lives

So far I have studied welfare at a fixed population. A different question is what ageing does to the utility of an individual. Consider a household born at date 0 on a balanced growth path. At age a she spends $\phi(a)E_a$, where average expenditure $E_a = (1 - \bar{v}\tau/\eta)\bar{Q}_a$ grows at the rate g . By (48), she consumes $C_a = \phi(a)E_a \mathcal{F}^h(a)^{1/(\eta-1)}$, and her expected lifetime utility (4) is

$$\mathcal{U}_0 = \int_0^\infty e^{-(\rho+\mu)a} \log C_a da.$$

As for social welfare in (53), lifetime utility has a closed form on a balanced growth path:

$$\begin{aligned}
(\rho + \mu) \mathcal{U}_0 = & \log \bar{Q}_0 + \underbrace{\frac{\log f_0}{\eta - 1}}_{\text{resources}} + \underbrace{\log \left(1 - \frac{\bar{v} \tau}{\eta}\right)}_{\text{life cycle}} + \underbrace{\log \left(1 - \frac{\phi}{\mu}\right) + \frac{\phi}{\rho + \mu}}_{\text{life cycle}} \\
& + \underbrace{\frac{1}{\eta - 1} \frac{g_f - (1 + g_q) \tau (1 - \mathcal{F}_{\rho + \mu}(0))}{\rho + \mu}}_{\text{familiarity}} + \underbrace{\frac{\tau g_q}{(\eta - 1)(\rho + \mu)}}_{\text{growth}}, \tag{55}
\end{aligned}$$

where $\mathcal{F}_{\rho + \mu}(0)$ is the index of a new product computed at the hazard $\rho + \mu$: the newborn weighs each age of her own life by impatience and survival, $(\rho + \mu)e^{-(\rho + \mu)a}$, rather than by the cross-section of the population.⁴²

Compared with social welfare (53), the four terms differ only in how they are discounted. The resource term is the same. Society values growth at g/ρ , while the newborn values it at $g/(\rho + \mu)$, since her death ends the flow. Society discounts the familiarity term only by death, at the rate μ , while she discounts it by death and impatience, at $\rho + \mu$. In the life-cycle term, she discounts the change of her own expenditure with age at $\rho + \mu$, while society averages it over the population. The death hazard thus also enters her utility through the discount. This is the gain from a longer life, which social welfare at a constant population does not count.

The effect of ageing on a newborn splits into a longevity effect, at a fixed balanced growth path, and the response of the equilibrium:

$$\frac{d\mathcal{U}_0}{d\mu} = \underbrace{\frac{\partial \mathcal{U}_0}{\partial \mu}}_{\text{longevity, path fixed}} + \underbrace{\frac{\partial \mathcal{U}_0}{\partial \tau} \cdot \frac{d\tau}{d\mu}}_{\text{equilibrium response}}.$$

This decomposition parallels the one for social welfare (54). At a given rate of creative destruction, ageing raises the newborn's lifetime utility, $\partial \mathcal{U}_0 / \partial \mu < 0$, as long as her flow utility $\log C_a$ is positive at every age. Under logarithmic utility, this is a normalisation of the units of consumption (with $\phi \neq 0$, it holds up to a term of second order in ϕ from the level of the income profile). Unlike for social welfare, the gain does not rest on familiarity: a longer life adds years of positive utility, so ageing would raise her utility at a given τ even if familiarity did not raise utility.

⁴²The familiarity term follows from integrating the household's law of motion in the proof of Lemma 3 against these weights. Impatience tilts them towards youth, so $\mathcal{F}_{\rho + \mu}(0) > \mathcal{F}_\mu(0)$: a new product commands a larger share of a young consumer's familiarity.

7.3 Quantitative effect of ageing on welfare

Table 8 reports the welfare effect of ageing in three exercises, at the calibration of Sections 4 and 5. The first compares the balanced growth paths of the 1980s and the 2010s. The other two follow the transitions of Section 5 and compare them with an economy that stays on the 1979 balanced growth path: in the first, ageing stops in 2025; in the second, it follows the Census Bureau projections to 2100. The total is the consumption equivalent (50). The other columns are the contributions of the four terms of (53), discounted and in log points. They add up to the total in logs, so the consumption equivalent slightly exceeds their sum. Figure 8 shows the same numbers, together with those for a newborn.

Table 8: The welfare effect of ageing, and where it comes from


	Resources	Life cycle	Growth	Familiarity	Total (CE)
Balanced growth paths	+2.03	+0.10	−3.00	+4.08	+3.26%
Transition, ageing stops in 2025	+0.68	+0.08	−0.95	+1.80	+1.62%
Transition, projections to 2100	+0.68	+0.08	−0.96	+1.95	+1.77%

Ageing raises welfare in all three exercises. Across balanced growth paths, the gain is 3.26% of consumption. Along the observed demographic path it is 1.62%, about half as large, because the entry boom of the 1980s delays the growth loss but also the familiarity gain. The projected ageing of the rest of the century raises the gain only to 1.77%, entirely through familiarity: the additional ageing comes late, and discounting mutes it, even though its effect on entry is sizeable. The life-cycle term is positive but small, about a tenth of a percentage point. Between the 1980s and the 2010s, the expenditure profile went from slightly falling with age to almost flat, so expenditure became more evenly spread across ages.⁴³

These numbers rest on the stance I took at the start of this section, that familiarity raises utility. Suppose instead that it does not: familiarity tilts demand towards the products a household knows, as in the rest of the paper, but adds nothing to their enjoyment. Then two things change. First, the familiarity term drops out of welfare. Second, the tilt becomes a cost: households buy too much of the products they know and too little of the others. With an aggregator that weights goods by quality alone, their consumption falls by $\log(\mathbb{E}_\omega[f^{(\eta-1)/\eta}]^{\eta/(\eta-1)}/\mathbb{E}_\omega[f]) \leq 0$, an average over the quality-weighted products the household buys. Older households, whose familiarity is

⁴³The positive effect of ageing on welfare depends on discounting. The growth loss compounds forever, while the gains in familiarity and resources are level effects. If I hold the equilibrium path fixed and lower the discount rate, the consumption equivalent of the observed path changes sign at $\rho^* \approx 0.021$. The calibrated $\rho = 0.05$ is above this threshold, so ageing raises welfare.

Figure 8: The welfare and utility effects of ageing, and where they come from



Consumption equivalents of ageing, in percent, and the contribution of each term, in percentage points (Tables 8 and 9). *Left*: social welfare (53). *Right*: the lifetime utility of a newborn (55). Balanced growth paths: the 1980s against the 2010s. Transition: the paths of Section 5 against the 1979 balanced growth path, with the mean adult age of spending frozen at 2025 or following the Census Bureau projections to 2100.

more uneven across products, misallocate more: across balanced growth paths, the average loss rises from 6.3 to 8.0 log points of consumption. The verdict then reverses. Without the familiarity term, ageing lowers welfare by 0.2% of consumption along the observed demographic path, in both scenarios, and by 0.9% across balanced growth paths, since the growth loss outweighs the resources saved on entry. With the misallocation, the losses are 0.9% along the observed path and 2.6% across balanced growth paths. Whether ageing raises or lowers welfare thus depends on whether familiarity is part of preferences. The positive results of the paper do not.

I also compute the effect of ageing on individual utility. Table 9 runs the same three exercises for the newborn of Section 7.2, evaluating (55) over her own life. The four terms of the social table reappear, together with longevity, which social welfare does not count.

Along the transition, the newborn gains far more than society: 6.2% of consumption against 1.6% along the observed path, and 6.5% against 1.8% with the projections. Almost all of her gain is longevity. The response of the equilibrium adds about 0.4 percentage points, made up of small resource, familiarity, growth and life-cycle terms. Across balanced growth paths, instead, her gain of 3.0% is slightly below society's 3.3%: longevity adds 4.7 points, but the life-cycle term subtracts 3.5. In the 1980s expenditure fell slightly with age, while in the 2010s it rises. A

Table 9: The welfare effect of ageing for a newborn

	Resources	Life cycle	Growth	Longevity	Familiarity	Total (CE)
Balanced growth paths	+2.15	−3.46	−1.88	+4.71	+1.40	+2.95%
Transition, ageing stops in 2025	+0.26	−0.12	−0.23	+5.60	+0.48	+6.17%
Transition, projections to 2100	+0.25	−0.10	−0.23	+5.93	+0.47	+6.52%

newborn of the old era therefore starts with less than the average household, and catches up only later in a life she discounts at $\rho + \mu$.

Three differences with Table 8 stand out. The growth term is small along the transition, -0.23 against -0.95 for society: the newborn is born into the entry boom, and the slowdown arrives late in her life, discounted at $\rho + \mu$ rather than ρ . The familiarity term is also small, $+0.48$ against $+1.80$, because the newborn starts at the reset level f_0 and builds her familiarity over her own life: she does not inherit the familiarity of the living population, which is what the social term counts. Finally, longevity contributes more along the transition than across balanced growth paths, $+5.6$ against $+4.7$, because the cohort born in 1980 lives through the whole observed fall in mortality: its expected adult lifespan is 30.3 years, against 25.1 on the 1979 balanced growth path.

8 Familiarity-destruction externality and the first best

Section 7 studied welfare along the equilibrium paths. This section asks which path a planner would choose instead, facing the same environment as the market economy. I obtain three results. First, at a given rate of creative destruction, the planner would produce, allocate and research as the market does, apart from how expenditure is spread across ages. This result relies on logarithmic utility, and implies that the only margin on which the equilibrium can fail is the rate of creative destruction. Second, the planner takes into account an effect that innovators ignore: a new product destroys the familiarity of consumers with the product it replaces. I call this effect familiarity destruction. It is a new externality, alongside the appropriability, intertemporal spillover and business-stealing effects of Aghion and Howitt (1992), and it pushes the decentralised economy towards too much creative destruction. In the calibrated economy, the planner values one more innovation at only 68% of its cost in the 1980s and 53% in the 2010s, and values even the first innovation at less than its cost, while the market sustains an entry rate above ten percent. Third, ageing also lowers the first-best rate of creative destruction, so the

slowdown of Section 4 is not, in itself, evidence of a market failure. As in Section 7, these results assume that familiarity raises utility: if it does not, the market innovates too little rather than too much.

The planner allocates labour to the production of goods, $y_j = l_j$, and goods to households, c_{ij} , and to research, x_j . It also chooses the rate of creative destruction $\tau_t \geq 0$. The planner uses the innovation technology of the market: innovations arrive on random lines, and the planner only chooses their rate.⁴⁴ Sustaining a rate τ requires a research aggregate $X \geq \nu\tau\bar{Q}_t$, and the planner maximises welfare (47).

The research basket has the same CES form as consumption. Producing the aggregate $\nu\tau\bar{Q}$ absorbs $\nu\tau$ units of labour, as I show in the proof of the next proposition.

Proposition 8. *Fix a date, the cross-section of states, and the innovation rate τ , hence $L^R = \nu\tau$. The planner's optimal allocation gives every household the same quantity of goods, $\int_0^1 c_{ij} dj = 1 - \bar{\nu}\tau/\eta$, allocated within the basket as $c_{ij} \propto q_j f_{ij}$; consumption at age a^h is*

$$C(a^h) = \left(1 - \frac{\bar{\nu}\tau}{\eta}\right) \bar{Q}_t \mathcal{F}_t^h(a^h)^{\frac{1}{\eta-1}}, \quad (56)$$

with $\mathcal{F}_t^h$ the average familiarity (24) at the date- t distributions; and when $\phi = 0$ it coincides with the market allocation at the same τ . Per-capita flow welfare is (49) with $E_t = (1 - \bar{\nu}\tau/\eta) \bar{Q}_t$ and no life-cycle term.

Logarithmic utility is key, as discussed in Section 3.2: familiarity shifts the log of consumption additively and leaves the marginal utility of goods unchanged. The planner therefore has no reason to favour the familiar old over the unfamiliar young. With an expenditure profile $\phi \neq 0$, the market's allocation differs from the planner's in one respect only: households of different ages have different budgets, $\phi(a^h)E$, while the planner gives everyone the same. This difference only changes the life-cycle term of welfare, which does not depend on τ . Uniform markups do not distort the static allocation either, since the wedge $1/\eta$ is common to every good. Every inefficiency that matters for innovation is therefore in τ .

The planner's effective discount rate is $\rho + (1 + g_q)\tau$: the market's discount in (33), $\bar{\rho} + (1 + g_q)\tau$, with the rate of time preference ρ in place of $\bar{\rho}$. To find the first-order condition, I add a small mass of innovations at one date to a constant path and follow its effects. Growth rises permanently by the quality step. Resources pay for the innovations at that date. Familiarity changes through a composition effect, which I derive in the proof of Proposition 9 in Appendix C: $(1 + g_q)[(\rho + (1 + g_q)\tau)\widehat{\mathcal{F}}(\rho + (1 + g_q)\tau) - 1]/\rho$ per unit of innovations. Setting the sum

⁴⁴A planner able to direct innovation across lines would trade off the quality gain against the familiarity stock destroyed on the chosen line, and would generally not choose uniform arrival; this is a richer technology than the one entrants operate.

of the three effects to zero gives the first-order condition. The planner averages familiarity over people, at the death hazard μ , while the market weights households by their expenditure, at the spending hazard $\tilde{\mu}$.

Proposition 9. (i) *An interior stationary first best τ^{FB} satisfies*

$$(1 + g_q) (\rho + (1 + g_q)\tau) \widehat{\mathcal{F}}(\rho + (1 + g_q)\tau; \mu, \tau) - 1 = \frac{(\eta - 1) \rho \bar{v}}{\eta - \bar{v} \tau}. \quad (57)$$

(ii) *At $g_f = 0$ the objective is strictly concave, the optimal path is constant, and*

$$\tau^{\text{FB}} \Big|_{g_f=0} = \max \left\{ \frac{\eta}{\bar{v}} - \frac{(\eta - 1) \rho}{g_q}, 0 \right\}. \quad (58)$$

(iii) *If*

$$(1 + g_q) \left[1 - \frac{\rho g_f}{(\mu + g_f)(\rho + \mu)} \right] < 1 + \frac{(\eta - 1) \rho \bar{v}}{\eta}, \quad (59)$$

no interior stationary first best exists near $\tau = 0$: the social value of the first innovation falls short of its cost, and stationary welfare decreases in τ throughout a neighbourhood of the corner.

The left-hand side of (57) uses the same transform that prices entry. Regrouping,

$$\underbrace{g_q}_{\text{growth}} + \underbrace{(1 + g_q) \rho \widehat{\mathcal{F}}(\rho + (1 + g_q)\tau)}_{\text{fresh line}} - \underbrace{(1 + g_q) \left[1 - (1 + g_q) \tau \widehat{\mathcal{F}}(\rho + (1 + g_q)\tau) \right]}_{\text{displaced pool}} = \underbrace{\frac{(\eta - 1) \rho \bar{v}}{\eta - \bar{v} \tau}}_{\text{resource cost}}. \quad (60)$$

An innovation brings the growth step and creates a new line. Consumers do not know the new line at first, and its familiarity index rises as they learn it. The line survives replacement at the rate $(1 + g_q)\tau$, and the planner discounts at ρ : survival and discounting combine into the rate $\rho + (1 + g_q)\tau$, so the new line is worth $\rho \widehat{\mathcal{F}}(\rho + (1 + g_q)\tau)$ per unit of time. But the innovation also destroys the familiarity of consumers with the line it replaces. Since $\mathcal{F}$ rises with age, creative destruction always replaces better-known products with less-known ones, so the two familiarity terms are negative on net. The profits lost by the displaced incumbent do not appear: they are a transfer to the entrant, the business-stealing wedge.

Multiplying both sides of free entry (33) by $(1 + g_q)(\eta - 1)\rho/(\eta - \bar{v}\tau)$ puts it in the same units:

$$\underbrace{\frac{\eta - 1}{\eta - \bar{v} \tau} (1 + g_q) \rho \widehat{\mathcal{F}}(\tilde{\rho} + (1 + g_q)\tau; \tilde{\mu}, \tau)}_{\text{entrant's share of the fresh line}} = \underbrace{\frac{(\eta - 1) \rho \bar{v}}{\eta - \bar{v} \tau}}_{\text{resource cost}}. \quad (61)$$

The market's condition (61) has the same resource cost as the planner's (60), but it values only the fresh line: it omits the growth step and the displaced pool. For an innovator, $\widehat{\mathcal{F}}$ is the dis-

counted familiarity path of its own product. For the planner, the new line's familiarity services follow the same path, while the familiarity destroyed on the replaced lines adds up, by the normalisation (23), to the constant 1. The expenditure profile adds two differences. The market weights households by their expenditure, at the spending hazard $\tilde{\mu}$, while the planner weights them equally, at the death hazard μ . And the market discounts at $\tilde{\rho}$, the planner at ρ . To see where the market and the planner differ, I compare their valuations of one more innovation at the market's rate of entry, margin by margin. Let $F_P \equiv (1 + g_q)\rho \widehat{\mathcal{F}}(\rho + (1 + g_q)\tau; \mu, \tau)$ be the planner's value of the new line and $F_M \equiv (1 + g_q)\rho \widehat{\mathcal{F}}(\tilde{\rho} + (1 + g_q)\tau; \tilde{\mu}, \tau)$ the market's, both in units of the flow of an average product line, and let $s \equiv (\eta - 1)/(\eta - \bar{\nu}\tau)$ be the share of its line's value that the entrant keeps. The planner's valuation minus the market's is

$$\underbrace{F_P - sF_M}_{\text{appropriability}} + \underbrace{g_q}_{\text{intertemporal spillover}} \underbrace{-F_P}_{\text{business stealing}} - \underbrace{(1 + g_q) \left[1 - (\rho + (1 + g_q)\tau) \widehat{\mathcal{F}}(\rho + (1 + g_q)\tau; \mu, \tau) \right]}_{\text{familiarity destruction}} . \quad (62)$$

Table 10 lists the four margins. The first three are the classic wedges of the Schumpeterian model. The entrant keeps only part of the value of its product, so appropriability makes the market innovate too little. The entrant also ignores that its quality step raises the quality of all future products, so the intertemporal spillover also makes the market innovate too little. On the other hand, the entrant values the sales it takes from the incumbent, which society does not, so business stealing makes the market innovate too much. The fourth wedge is new. The entrant ignores the familiarity that consumers had built with the product it replaces, so familiarity destruction makes the market innovate too much.

Table 10: The value of one more innovation: margins and wedges

Margin	Wedge	Planner vs market
Sales of the new product over its life	Appropriability	+
Growth step	Intertemporal spillover	+
Sales taken from the replaced product	Business stealing	−
Familiarity with the replaced product	Familiarity destruction	−
Research cost	None	=

Notes. The sign of each wedge of (62): + when the planner values the margin more than the market, so that the market innovates too little; − when it values it less.

Figure 9 reports the wedges at the calibration of Section 4, as shares of the research cost. In the 1980s, the classic wedges add up to +47%: on their own, they call for more innovation. Familiarity destruction subtracts 79%, so the planner values one more innovation at only 68% of

its cost. In the 2010s, the classic wedges add up to +50% and familiarity destruction subtracts 97%: the planner values one more innovation at 53% of its cost. Ageing deepens familiarity destruction, because an older population has more familiarity for each innovation to destroy.


These wedges, like the welfare gains of Section 7, rest on familiarity raising utility. Suppose instead that familiarity is only a friction that tilts demand towards the products households know. Then two things change. First, the static allocation is no longer efficient: households spend too much on familiar products, and the planner would allocate each basket by quality alone. Second, familiarity destruction turns from a cost into a gain, since resetting familiarity reduces this misallocation.⁴⁵ In the calibrated economy, the familiarity wedge becomes +19% of the research cost in the 1980s and +25% in the 2010s, instead of −79% and −97%. With the classic wedges, the planner would then value one more innovation at 166% of its cost in the 1980s and 175% in the 2010s: the market would innovate too little, and more so in the older economy. Whether the market innovates too much or too little thus depends on whether familiarity is part of preferences.

Familiarity destruction also rests on a second assumption: a new product drives the old one out of its line (Section 3.4). In reality, old products need not exit. A firm that launches a new product can keep selling the old one, and households that know the old product can keep buying it. Their familiarity then survives the innovation. Many old products do exit, however. Each product carries fixed costs of production and distribution, and in a longer product line these costs can outweigh the extra sales (Bayus and Putsis, 1999). Retailers delist products whose sales fall, and firms' own new products hasten the decline of their old ones (Argente et al., 2024). Consumers miss the products that disappear. When a firm eliminates a product, its customers' satisfaction and loyalty fall (Homburg et al., 2010); when a retailer delists a brand or cuts its assortment, some customers switch store or buy less (Borle et al., 2005; Sloot and Verhoef, 2008); and consumers campaign to bring discontinued brands back (Davari et al., 2017; Almazyad et al., 2023). When Coca-Cola replaced its formula with New Coke in 1985, the protests led it to bring back the old formula as Coca-Cola Classic within three months. The firm that withdraws a product weighs the profits the product would still earn, not the surplus of the consumers who know it (Spence, 1976), so familiarity destruction is present even when exit is a choice. The size of the wedge, however, should not be taken literally. It is computed as if every replaced product left the market, and it overstates familiarity destruction to the extent that replaced products

⁴⁵Under this view, the familiarity term of welfare is replaced by the misallocation term of Section 7.3, $\mathbb{E}[\frac{\eta}{\eta-1} \log(\mathcal{F}_c^h(a^h)/f_0) - \log(\mathcal{F}^h(a^h)/f_0)]$, where $\mathcal{F}_c^h$ is the average familiarity (24) computed with familiarity growth $c = (\eta - 1)g_f/\eta$ in place of g_f . Applying the composition effect of Proposition 9 to each of the two averages, the familiarity-destruction wedge of (62) becomes $(1 + g_q)[\eta((\rho + (1 + g_q)\tau)\widehat{\mathcal{F}}_c - 1) - (\eta - 1)((\rho + (1 + g_q)\tau)\widehat{\mathcal{F}} - 1)]$, where both transforms are evaluated at $\rho + (1 + g_q)\tau$ and $\widehat{\mathcal{F}}_c$ is the transform of the index computed with familiarity growth c . The two terms cancel to first order in g_f : misallocation is a second-order effect of the dispersion of familiarity across products.

survive.

Figure 9: The value of one more innovation: planner minus market



The wedges of (62) at the market's rate of entry, in percent of the research cost of the innovation, on the balanced growth paths of the 1980s and the 2010s (Section 4). The market values the marginal innovation at its cost, the planner at its cost plus the net wedge. The wedges assume that familiarity raises utility: if it does not, familiarity destruction becomes a gain of 19% of the research cost in the 1980s and 25% in the 2010s, and the net wedge is +66% and +75%.

The main comparative static of the paper also holds for the planner.

Proposition 10. Fix $0 < \mu_1 < \mu_2$. There exists $\check{g}_f > 0$ such that, for every $g_f \in (0, \check{g}_f]$ and $\mu \in [\mu_1, \mu_2]$, any interior stationary first best with $g_f < (1 + g_q) \tau^{\text{FB}}$ responds to the death hazard by

$$\frac{d\tau^{\text{FB}}}{d\mu} = \frac{(1 + g_q)(\rho + (1 + g_q)\tau) \partial_{\mu} \widehat{\mathcal{F}}}{\frac{(\eta - 1) \rho \bar{v}^2}{(\eta - \bar{v} \tau)^2} - \Phi_{\tau}} > 0, \quad (63)$$

where $\Phi(\tau) \equiv (1 + g_q)(\rho + (1 + g_q)\tau) \widehat{\mathcal{F}}(\rho + (1 + g_q)\tau; \mu, \tau)$ is the first term of (57). Ageing lowers the first-best rate of creative destruction, and with it first-best growth $g_q \tau^{\text{FB}} / (\eta - 1)$.

The threshold $\check{g}_f$ is needed only to sign the denominator, by the same continuity argument as in Proposition 4: at $g_f = 0$, $\Phi \equiv 1 + g_q$ is flat, and it stays nearly flat for small g_f .

An older population is worse at absorbing novelty. This is a feature of preferences and demography, not of market structure, so the planner, too, responds to ageing by creating fewer new products. The slowdown of Section 4 is therefore not, in itself, a sign of a market failure. Whether the gap between the market and the first best widens as μ falls is a quantitative question, and Figure 9 answers it. At the market's rate of entry, the planner's net valuation of one more innovation falls from -32% to -47% of its cost: the market's excess of innovation over the first best grows with ageing.

9 Model predictions in the data

The model was calibrated to entry and growth in the 1980s and to how the share of spending on new products changes with the age of households, and used to shed light on how much of the growth slowdown ageing explains. However, it has further implications that the calibration does not target. This section states four of them as predictions and takes each to the data. The predictions analysed concern how firms launch products and the size of firms in markets in which buyers are of different ages, how the value of firms evolves with firm age, and how consumers' ageing changes the firm value profile. Before these, I look at the trend in research and development: as the population ages, consumer-facing industries should spend less on it.

First, though, I need to be explicit about what I mean by markets where consumers have different ages. The model has one single market, while the data contain many: industries differ in who buys from them. To bring the second and third predictions to these data, I suppose that the economy is made of many markets, each a copy of the economy of Section 3, with its own product lines and its own entrants, who pay the same research cost for innovating. Markets differ only in the age of their buyers: the dollars spent in market m have the mean adult age $1/\tilde{\mu}_m$. Each market has its own balanced growth path, and comparing markets amounts to comparing balanced growth paths that differ in $\tilde{\mu}$.

9.1 The trend in R&D spending in packaged-goods industries


Prediction 1. As the population grows older, spending on research and development decreases in consumer-facing industries.

In the model, research spending equals the cost of entry: entrants pay $v\bar{Q}$ per innovation, so a market spends $v\tau\bar{Q}$ on research, a share of sales that moves with the entry rate. Where buyers

are older, a new product is worth less, fewer entrants pay the research cost, and research per dollar of sales is lower.

To check this prediction in the data, I use the National Science Foundation's surveys of business R&D, which cover about twenty industries that sell a large part of their output to consumers.

Figure 10: R&D spending by consumer packaged-goods industries, 2013–2023



R&D performed by food (NAICS 311), beverages and tobacco (312), apparel (313–316) and soap and toiletries (3256), billions of 2023 dollars; dashed: least-squares trend, -0.07 billion a year. Sources: NSF business R&D surveys; BEA.

Figure 10 plots the R&D expenditure of the industries that make the products captured by the HMS: food, beverages and tobacco, apparel, and soap and toiletries.⁴⁶ Over time, it fell from \$11.6 billion in 2013 to \$10.5 billion in 2023, 9% in 2023 dollars, while the R&D of all US businesses rose by 73%, from \$418 to \$722 billion; over the same years the mean age of the buyers of food and beverages rose from 49.6 to 51.0 years. The industries whose products are affected by older people's inertia cut research, even while the rest of the economy raised it, as the model predicts.

9.2 Fewer new products in markets where buyers are older

Prediction 2. Markets whose buyers are older have a lower rate of entry of new products.

The prediction is Proposition 5 applied market by market: older buyers can mean both that consumers are older on average and that old consumers spend more than younger ones. In each

⁴⁶Domestic R&D performed by companies with ten or more employees, NSF Business R&D and Innovation Survey (2013–2016) and Business Enterprise R&D Survey (2017–2023), deflated with the GDP price index; earlier years do not report some of these industries.

market, free entry equates the value of a new product to the cost of research. Where buyers are older, the net present value of profits from a new product is lower, so fewer products are launched.

The HMS allows a test of this prediction, since product modules differ in the age of the households that buy them. For each module I measure the spending-weighted mean age of the heads of the households with positive expenditure on it, and its entry rate in each year, the number of new products launched over the number of products on sale the year before. Since the entry rate influences the contemporaneous age of a module's buyers (as I documented, young households spend more on new products), I use a predetermined measure of the age of a module's buyers. I perform two exercises: in the first, I measure the age of a module's buyers once, in 2004, the first year of the panel and the year before the first launches; in the second, I use the module's exposure to young buyers in the previous year, built like the exposure to new goods of Section 2: the age its buyers would have if its 2004–2006 customers had followed the spending of their age groups in the rest of the panel.⁴⁷

Figure 11 plots the entry rate against the two measures, net of year and product-group effects; Table 11 reports the Poisson regressions. Within product groups, the entry rate falls with the age of a module's dollars: a module whose buyers were ten years older in 2004 launches 20% fewer new products per product on sale (standard error 7%), and 14% fewer once I also control for how old the module's brands already were in 2004, the spending-weighted mean age of its brands dated by the first use of their trademarks. With the exposure to young buyers in the previous year, the fall is 23% (standard error 7%), and 15% with the age of the module's brands also measured in the previous year. The point estimates are higher than what the model predicts: in the calibration to the 2010s, the entry rate is about 6.5% lower per ten years of the mean age of a market's dollars.⁴⁸ However, with the controls for maturity the gradient is not significantly steeper than the model's.


9.3 Young firms are smaller where buyers are older

In the model, each product line is run by one firm, and the age of the firm is the age of its product. By the demand (26), at the common price the sales of a firm are proportional to $q_j \mathcal{F}(a_j^p)$. By Lemma 2 and Lemma 1, the expected quality of a line of age a^p is $(1 + g_q)e^{-g_q \tau a^p} \bar{Q}^{\eta-1}$, so

⁴⁷ Shares of the module's 2004–2006 dollars by five-year age group of the head, each scaled by the growth of that group's spending on other product groups; the exposure is the mean age of the predicted dollars. It moves with the population, not with the module's launches, and tracks the actual age of the previous year's buyers closely (coefficient 0.87, standard error 0.02).

⁴⁸ I calibrate the model to 2010–2019 as in Section 4.2: entry of 8.16% and growth of 1.64% give $g_q = 0.401$, the age profile of spending on new products gives $g_f = 0.064$, and free entry gives $\bar{v} = 7.29$. Markets whose dollars are five years younger than the 2010s mean of 29.5 years above 18 have an entry rate of 8.43%, five years older 7.89%; the calibration to the 1980s gives 6.3%.

Figure 11: The entry rate of new products and the age of a market's buyers



The figure shows the entry rate of new products in a product module, launches over the products on sale the year before, in logs, against the mean age of its buyers in 2004 in panel (a) and its exposure to young buyers in the previous year, the predicted mean age of its buyers, in panel (b), both net of year and product-group effects; NielsenIQ Homescan, 2005–2024. Dark-blue dots are twenty equal-size bins of module-years, and the dark-blue line the least-squares fit, with its slope per ten years of buyer age and standard error, clustered by module, in the legend. Light-blue diamonds and line add a control for the age of the module's brands. The dashed copper line is Prediction 2 in the model calibrated to the 2010s, -6.5% per ten years. The regressions are in Table 11. Sources: NielsenIQ Homescan; USPTO.

Table 11: The entry rate of new products and the age of a market's buyers

	Buyers in 2004			Exposure in the previous year		
	(1)	(2)	(3)	(4)	(5)	(6)
Mean age of buyers (per ten years)	-0.008 (0.229)	-0.225*** (0.068)	-0.156* (0.067)	-0.036 (0.216)	-0.264*** (0.066)	-0.161* (0.068)
Dollars from heads 65 and over (per 10 pp)	-0.072 (0.115)	-0.149*** (0.035)	-0.118*** (0.036)	-0.066 (0.101)	-0.144*** (0.034)	-0.096** (0.036)
Dollars from heads under 35 (per 10 pp)	-0.056 (0.138)	+0.106** (0.034)	+0.068* (0.032)	-0.020 (0.122)	+0.127*** (0.034)	+0.065 (0.034)
Year effects	Yes	Yes	Yes	Yes	Yes	Yes
Product-group effects	No	Yes	Yes	No	Yes	Yes
Age of the module's brands in 2004	No	No	Yes	No	No	No
Age of the module's brands in the previous year	No	No	No	No	No	Yes
Modules	1,015	1,015	1,013	1,015	1,015	1,014
Module-years	20,216	20,216	20,188	20,216	20,216	20,148
New products per 100 products on sale		4.55			4.55	

Notes. Poisson regressions of the number of new products launched in a module in a year on the age of its buyers, with the log of the products on sale the year before as a regressor; a coefficient of -0.1 means about 10% fewer launches. One regressor per row, in ten-year or ten-percentage-point units, measured on the module's buyers in 2004 in columns (1) to (3) and predicted from the module's exposure to young buyers in the previous year in columns (4) to (6). Columns (3) and (6) add the age of the module's brands, in 2004 and in the previous year. 2005–2024. Standard errors clustered by module in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Sources: as in Figure 11.

that the expected sales of a firm of age zero relative to a firm of age a^p are

$$\frac{\mathbb{E}[y_j \mid a_j^p = 0]}{\mathbb{E}[y_j \mid a_j^p = a^p]} = e^{g_q \tau a^p} \frac{\mathcal{F}(0)}{\mathcal{F}(a^p)}. \quad (64)$$

A young firm has the better product, but its buyers are not yet familiar with it, and the penalty is stronger the older the consumers are. It can be shown formally that, given τ , the ratio (64) is increasing in $\tilde{\mu}$. In addition, when the entry rate of each market adjusts as in Prediction 2, so that older consumers also cause lower entry, this slower entry also shrinks the quality lead of young firms in older markets, and the gap between entrants and older products widens further.

Prediction 3. Given the entry rate, young firms are smaller relative to old firms in markets whose buyers are older.


I use Compustat to analyse sales of listed firms in consumer industries and find that listed firms confirm Prediction 3. For 47 consumer industry groups I measure the mean age of their buyers from the Consumer Expenditure Survey, as the spending-weighted mean age of the household heads who buy what the industry sells: 36 years for child care, 42 for homebuilders, 46 for apparel and restaurants, 50 for cable television, 54 for drug stores and casinos.⁴⁹ Figure 12 plots, for each group, the difference between the log sales of young firms, in their first five years in Compustat, and of mature firms, present for more than ten, against the age of the group's buyers, and Table 12 reports the regressions. As predicted by the model, younger firms are smaller compared to older ones in markets where buyers are older, and they are so by 0.64 log points per ten years of buyers' age in the regression with group-by-year effects. The estimates of Table 12 are above the model's predictions, but not precisely enough to say by how much, and with standard errors clustered by industry group they do not reject the model's value.

9.4 The price-earnings ratio falls with firm age and ageing makes it flatter

The third prediction concerns the value of firms. By Proposition 2, the value of a firm is its profits times the multiple $\kappa(a^p)$, a price-earnings ratio.

⁴⁹Compustat, US firms in the 47 industry groups, 1990–2024, 27,375 firm-years of 2,513 firms; groups are sets of four-digit SIC industries matched to the items of the Consumer Expenditure Survey, interview and diary. The age of a group's buyers in a fiscal year is the mean age of the heads of the households that bought its items in the surveys of that year and of the two adjacent years, each household weighted by its spending on them; the three years are pooled because a single survey year has few buyers of some items. Across groups, this age has a standard deviation of three years. Young firms are those in their first five years in Compustat, mature firms those present for more than ten; the regression is of log sales on a young-firm indicator and its interaction with the group's buyer age, with group-by-year effects; standard errors are clustered by firm. The figure plots the group means of the gap, demeaned by year.

Figure 12: The size of young firms relative to mature firms and the age of the industry's buyers



The figure shows the size of young firms relative to mature firms in an industry against the age of the industry's buyers; Compustat, 47 consumer industry groups, 1990–2024. Each point is a group: the mean log sales of firms in their first five years in Compustat minus that of firms present for more than ten, demeaned by year, against the mean age of the group's buyers in the Consumer Expenditure Survey; the area of a point is the number of firms. The solid line is the least-squares fit weighted by the number of firms, with its slope per ten years and robust standard error in the legend; the dashed copper line is Prediction 3 in the model calibrated to the 2010s, -0.13 per ten years. The regressions are in Table 12. Sources: Compustat; Consumer Expenditure Survey.

Table 12: The size of young firms relative to mature firms and the age of the industry's buyers: regressions

Young firms:	Log sales	
	first two years (1)	first five years (2)
Young firm	−1.805*** (0.102)	−1.537*** (0.097)
Young firm × mean age of buyers (per ten years)	−0.746** (0.282)	−0.643* (0.265)
Industry group × year effects	Yes	Yes
Firm-years	18,678	22,358
Firms	2,356	2,464
Industry groups	47	47

Notes. Regressions of log sales on an indicator for young firms and on its interaction with the mean age of the industry group's buyers in the fiscal year, in ten-year units and centred at 48 years. Young firms are those in their first two years in Compustat in column (1) and in their first five years in column (2); mature firms are those present for more than ten years, and the years in between are excluded. Industry-group-by-year effects, so that young firms are compared with the mature firms of their group and year. Standard errors clustered by firm in parentheses; by industry group, 0.37 and 0.36 for the interaction. The model's interaction is -0.13 . * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Sources: as in Figure 12.

Prediction 4. Under the conditions of Proposition 2: (i) the ratio of the value of a firm to its profits falls with the age of the firm after launch; (ii) given the entry rate, older buyers raise this ratio at every age and flatten its fall at launch.

Neither part of the prediction arises in standard quality-ladder models. Without familiarity growth, $g_f = 0$, the familiarity index is flat and $\kappa = 1/(\tilde{\rho} + (1 + g_q)\tau)$ at every age, as in Aghion and Howitt (1992) and in Klette and Kortum (2004): the value of a product line is a constant multiple of its profits, the same whether the line was created last year or thirty years ago, and the age of buyers has nothing to act on.


The value-profits decline can be documented in the data. Figure 13 plots the ratio of market value to earnings of listed US firms by the number of years since the firm was listed, net of industry and year effects.⁵⁰ The multiple falls by 0.23 log points between the first three years after listing and years three to five, and by 0.44 by years 31 to 40; among consumer-facing firms by 0.32, and then stays about 0.36 below its initial level. In the model the fall is 0.11 by age ten and 0.14 by age forty. That multiples fall with the age of the firm is known (Pástor and Veronesi, 2003; Loderer et al., 2017), and models in which demand accumulates over the life of a firm share the sign (Gourio and Rudanko, 2014; Foster et al., 2016). What the present model adds is part (ii), the link to the age of buyers.

The relationship between the flattening of the ratio and population ageing can be tested across decades, since the buyers of US firms aged between the 1980s and the 2010s. The bottom panels of Figure 13 estimate the profile decade by decade and find that the profile flattens. A firm listed 31 to 40 years earlier trades 0.59 log points below a newly listed firm in the 1980s, 0.54 in the 1990s, 0.45 in the 2000s and 0.24 in the 2010s. Among consumer-facing firms the profiles are flatter and less precisely estimated. Neither the number of firms listed in the same year as the firm nor the firm's own size at listing, each interacted with age, moves these slopes, and neither does the real interest rate.⁵¹ The flattening is in the direction the model predicts and about three times its size.

⁵⁰Compustat, US non-financial firms, fiscal years 1990–2024, 104,523 firm-years of 10,469 firms. The multiple is market equity at the end of the fiscal year over operating income before depreciation, for firm-years with positive operating income, winsorised at the first and last percentiles within each year; years since listing come from the first CRSP link of the firm and are top-coded at forty; the effects are two-digit SIC industry by fiscal year; standard errors are clustered by firm. Consumer-facing firms are the 2,199 firms in the 47 consumer industry groups of Section 9.3, with industry-group by year effects. The profile is the same when age is counted from the firm's first year in Compustat rather than from its listing.

⁵¹In a pooled regression over 1980–2019 with the slope on log age allowed to differ by decade, adding the interaction of log age with the real ten-year Treasury rate (the nominal yield less CPI inflation) leaves the decade slopes at -0.19 , -0.18 , -0.12 and -0.08 , and the interaction itself at -0.001 (standard error 0.002). Interacting log age with the log number of listings in the firm's listing year changes nothing; interacting it with the firm's size at listing, which is only observed for firms listed from 1980, steepens the profile in every decade by about the same amount and leaves the flattening, from -0.28 in the 1990s to -0.09 in the 2010s. Fiscal years 2020–2024 are left out.

Figure 13: Valuation multiples by firm age



The figure shows the age profile of valuation multiples. Top: the log of market equity over operating income before depreciation by years since listing, relative to the first three years, from a regression with industry-by-year effects; shaded areas are 95 per cent confidence bands, standard errors clustered by firm. All US non-financial firms (10,469) and consumer-facing firms (2,199), 1990–2024. Bottom: the same profile estimated by decade, 1980–2019, for all non-financial firms (left; 6,539, 7,494, 5,491 and 3,814 firms by decade) and for consumer-facing firms (right; 1,619, 1,719, 1,157 and 743); vertical bars are 90 per cent confidence intervals. Sources: Compustat; CRSP.

10 Conclusion

Using 111,002 new product launches and the purchases of about 50,000 US households a year, I documented that older households spend less on new products. Relative to households aged 30, the share of their spending on new products is 15% lower at age 50, 29% lower at 65 and 35% lower at 80. The gap is robust to controls for income, education, household composition, shopping behaviour and the categories households buy. When coupled with population ageing, this fact implies that in older societies new products sell less, and that ageing harms incentives to innovate. To study this argument, I developed a Schumpeterian model in which consumers become familiar with products over time. Older consumers have had more time to become familiar with existing products, so they spend less on new ones. As the population ages and spending shifts towards older households, new products sell less, the value of bringing them to the market falls, and entry and growth decline. In the model calibrated to the United States, the ageing between the 1980s and the 2010s explains 32.1% of the growth slowdown. This demand channel complements the supply channels through which demographics affect innovation and growth.

These effects take time to unfold. Anticipated ageing first raises growth: an innovator who enters before the population has aged sells at launch as much as before, but expects to sell more later, once consumers have aged. Then growth falls, and it keeps falling after the population stops ageing. Had the US kept its 1980 demographics, its GDP per capita would have been 3.5% higher in 2025, and, even if ageing stopped in 2025, the gap would reach 8.0% by 2050. The Census Bureau, however, projects that ageing will continue to the end of the century.

Ageing also changes the effectiveness of policy instruments: a subsidy of 10% of the cost of research raises growth by 11.0% more in the economy of the 2010s than in that of the 1980s. At the same time, slower growth need not mean lower welfare. Longer lives give consumers more time to become familiar with products, and slower creative destruction destroys less of their familiarity. If familiarity raises utility, ageing raises welfare by 3.3% in consumption equivalents across balanced growth paths. The model also uncovers a new externality: innovators ignore the familiarity that consumers have built with the products they replace.

The familiarity mechanism raises further questions. Firms can advertise rather than wait for familiarity to build, and in an older society they may shift resources from improving products to promoting them. In the data, unlike in the model, most spending on new products goes to established firms, which extend brands that consumers already know; if extensions carry familiarity over, incumbents have an advantage over entrants that grows with the age of consumers. Finally, process innovation lowers the cost of old products without resetting familiarity, and ageing may tilt innovation towards it. I plan to study these questions next.

References

- Acemoglu, D., D. Autor, K. Beirne, and A. Scott (2026). Baby busts and growth booms: Demographic change and the macroeconomy. NBER Working Paper 35401, National Bureau of Economic Research.
- Acemoglu, D. and J. Linn (2004). Market size in innovation: Theory and evidence from the pharmaceutical industry. *Quarterly Journal of Economics* 119(3), 1049–1090.
- Acemoglu, D. and P. Restrepo (2017). Secular stagnation? The effect of aging on economic growth in the age of automation. *American Economic Review* 107(5), 174–179.
- Acemoglu, D. and P. Restrepo (2022). Demographics and automation. *The Review of Economic Studies* 89(1), 1–44.
- Aghion, P., A. Bergeaud, T. Boppart, P. J. Klenow, and H. Li (2023). A theory of falling growth and rising rents. *Review of Economic Studies* 90(6), 2675–2702.
- Aghion, P. and P. Howitt (1992). A model of growth through creative destruction. *Econometrica* 60(2), 323–351.
- Aguiar, M. and E. Hurst (2007). Life-cycle prices and production. *American Economic Review* 97(5), 1533–1559.
- Akcigit, U. and S. T. Ates (2021). Ten facts on declining business dynamism and lessons from endogenous growth theory. *American Economic Journal: Macroeconomics* 13(1), 257–298.
- Akcigit, U. and S. T. Ates (2023). What happened to US business dynamism? *Journal of Political Economy* 131(8), 2059–2124.
- Almazyad, F., P. Shah, and E. T. Loiacono (2023). Social media activism for resurrecting deleted brands: The role of consumers' psychological reactance. *Journal of Brand Management* 30(4), 367–380.
- Argente, D., S. Baslandze, D. Hanley, and S. Moreira (2025). Patents to products: Product innovation and firm dynamics. NBER Working Paper 34592, National Bureau of Economic Research.
- Argente, D., M. Lee, and S. Moreira (2024). The life cycle of products: Evidence and implications. *Journal of Political Economy* 132(2), 337–390.
- Arts, J. W. C., R. T. Frambach, and T. H. A. Bijmolt (2011). Generalizations on consumer innovation adoption: A meta-analysis on drivers of intention and behavior. *International Journal of Research in Marketing* 28(2), 134–144.

- Auclert, A., B. Bardóczy, M. Rognlie, and L. Straub (2021). Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models. *Econometrica* 89(5), 2375–2408.
- Auclert, A., E. Majic, M. Rognlie, and L. Straub (2025). Determinacy and large-scale solutions in the sequence space. *Journal of Political Economy Macroeconomics*. Forthcoming.
- Basker, E. (2012). Raising the barcode scanner: Technology and productivity in the retail sector. *American Economic Journal: Applied Economics* 4(3), 1–27.
- Basker, E. and T. Simcoe (2021). Upstream, downstream: Diffusion and impacts of the universal product code. *Journal of Political Economy* 129(4), 1252–1286.
- Bayus, B. L. and W. P. Putsis (1999). Product proliferation: An empirical analysis of product line determinants and market outcomes. *Marketing Science* 18(2), 137–153.
- Blanchard, O. J. (1985). Debt, deficits, and finite horizons. *Journal of Political Economy* 93(2), 223–247.
- Blanchard, O. J. and C. M. Kahn (1980). The solution of linear difference models under rational expectations. *Econometrica* 48(5), 1305–1311.
- Bloom, N., C. I. Jones, J. Van Reenen, and M. Webb (2020). Are ideas getting harder to find? *American Economic Review* 110(4), 1104–1144.
- Borle, S., P. Boatwright, J. B. Kadane, J. C. Nunes, and G. Shmueli (2005). The effect of product assortment changes on customer retention. *Marketing Science* 24(4), 616–622.
- Bornstein, G. (2026). Entry and profits in an aging economy: The role of consumer inertia. *Review of Economic Studies*. Accepted; NBER Working Paper 33820.
- Bornstein, R. F. (1989). Exposure and affect: Overview and meta-analysis of research, 1968–1987. *Psychological Bulletin* 106(2), 265–289.
- Brès, M. and D. Angelini (2025). Aging consumers, competition, and growth. Working paper.
- Broda, C. and D. E. Weinstein (2006). Globalization and the gains from variety. *Quarterly Journal of Economics* 121(2), 541–585.
- Broda, C. and D. E. Weinstein (2010). Product creation and destruction: Evidence and price implications. *American Economic Review* 100(3), 691–723.
- Bronnenberg, B. J., J.-P. H. Dubé, and M. Gentzkow (2012). The evolution of brand preferences: Evidence from consumer migration. *American Economic Review* 102(6), 2472–2508.


- Cravino, J., A. A. Levchenko, and S. Rojas (2022, October). Population aging and structural transformation. *American Economic Journal: Macroeconomics* 14(4), 68–98.
- Crouzet, N., P. Ghosh, A. Gupta, and F. Mezzanotti (2026). Demographics and technology diffusion: Evidence from mobile payments. Working paper, Northwestern University.
- Dang, D. (2025). Customer capital and firm innovation. Working paper, University of Wisconsin–Madison.
- Davari, A., P. Iyer, and F. Guzmán (2017). Determinants of brand resurrection movements: Why consumers want dead brands back? *European Journal of Marketing* 51(11/12), 1896–1917.
- De Ridder, M. (2024). Market power and innovation in the intangible economy. *American Economic Review* 114(1), 199–251.
- Deaton, A. (1997). *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*. Baltimore: Johns Hopkins University Press for the World Bank.
- Decker, R., J. Haltiwanger, R. Jarmin, and J. Miranda (2014). The role of entrepreneurship in US job creation and economic dynamism. *Journal of Economic Perspectives* 28(3), 3–24.
- Einiö, E., J. Feng, and X. Jaravel (2025). Social push and the direction of innovation. Working paper.
- Engbom, N. (2019, April). Firm and worker dynamics in an aging labor market. Working Paper 756, Federal Reserve Bank of Minneapolis.
- Fernald, J. G. (2014). Productivity and potential output before, during, and after the great recession. *NBER Macroeconomics Annual* 29, 1–51.
- Foster, L., J. Haltiwanger, and C. Syverson (2016). The slow growth of new plants: Learning about demand? *Economica* 83(329), 91–129.
- Gielens, K. and J.-B. E. M. Steenkamp (2007). Drivers of consumer acceptance of new packaged goods: An investigation across products and countries. *International Journal of Research in Marketing* 24, 97–111.
- Gokhale, J., L. J. Kotlikoff, and J. Sabelhaus (1996). Understanding the postwar decline in U.S. saving: A cohort analysis. *Brookings Papers on Economic Activity* 27(1), 315–407.
- Gourio, F. and L. Rudanko (2014). Customer capital. *Review of Economic Studies* 81(3), 1102–1136.

- Homburg, C., A. Fürst, and J.-K. Prigge (2010). A customer perspective on product eliminations: How the removal of products affects customers and business relationships. *Journal of the Academy of Marketing Science* 38(5), 531–549.
- Hopenhayn, H., J. Neira, and R. Singhanian (2022). From population growth to firm demographics: Implications for concentration, entrepreneurship and the labor share. *Econometrica* 90(4), 1879–1914.
- Ignaszak, M. and P. Sedláček (2026). Customer acquisition, business dynamism and aggregate growth. *Review of Economic Studies*. Forthcoming.
- Jaravel, X. (2019). The unequal gains from product innovations: Evidence from the U.S. retail sector. *Quarterly Journal of Economics* 134(2), 715–783.
- Karahan, F., B. Pugsley, and A. Şahin (2024). Demographic origins of the start-up deficit. *American Economic Review* 114(7), 1986–2023.
- Klette, T. J. and S. Kortum (2004). Innovating firms and aggregate innovation. *Journal of Political Economy* 112(5), 986–1018.
- Krein, M. G. (1962). Integral equations on a half-line with kernel depending upon the difference of the arguments. *American Mathematical Society Translations, Series 2* 22, 163–288.
- Lambert-Pandraud, R. and G. Laurent (2010). Why do older consumers buy older brands? The role of attachment and declining innovativeness. *Journal of Marketing* 74(5), 104–121.
- Lambert-Pandraud, R., G. Laurent, and E. Lapersonne (2005). Repeat purchasing of new automobiles by older consumers: Empirical evidence and interpretations. *Journal of Marketing* 69(2), 97–113.
- Lehr, N. H. (2026). The demographics of technology adoption and productivity growth. Working paper, International Monetary Fund.
- Liu, E., A. Mian, and A. Sufi (2022). Low interest rates, market power, and productivity growth. *Econometrica* 90(1), 193–221.
- Loderer, C., R. Stulz, and U. Waelchli (2017). Firm rigidities and the decline in growth opportunities. *Management Science* 63(9), 3000–3020.
- Maestas, N., K. J. Mullen, and D. Powell (2023). The effect of population aging on economic growth, the labor force, and productivity. *American Economic Journal: Macroeconomics* 15(2), 306–332.

- Michelacci, C., L. Paciello, and A. Pozzi (2022). The extensive margin of aggregate consumption demand. *Review of Economic Studies* 89(2), 909–947.
- Neiman, B. and J. Vavra (2023). The rise of niche consumption. *American Economic Journal: Macroeconomics* 15(3), 224–264.
- Olmstead-Rumsey, J. (2026). Market concentration and the productivity slowdown. Working paper.
- Pástor, L. and P. Veronesi (2003). Stock valuation and learning about profitability. *Journal of Finance* 58(5), 1749–1789.
- Peters, M. and C. Walsh (2026). Population growth and firm dynamics. *Journal of Political Economy Macroeconomics* 4(2), 256–296.
- Rachel, L. (2024). Leisure-enhancing technological change. Working paper.
- Ravn, M. O., S. Schmitt-Grohé, and M. Uribe (2006). Deep habits. *Review of Economic Studies* 73(1), 195–218.
- Rogers, E. M. (1958). Categorizing the adopters of agricultural practices. *Rural Sociology* 23, 345–354.
- Rogers, E. M. (1983). *Diffusion of Innovations* (Third ed.). New York: The Free Press.
- Ryan, B. and N. C. Gross (1943). The diffusion of hybrid seed corn in two Iowa communities. *Rural Sociology* 8(1), 15–24.
- Sloot, L. M. and P. C. Verhoef (2008). The impact of brand delisting on store switching and brand switching intentions. *Journal of Retailing* 84(3), 281–296.
- Spence, M. (1976). Product selection, fixed costs, and monopolistic competition. *Review of Economic Studies* 43(2), 217–235.
- Steenkamp, J.-B. E. M. and K. Gielens (2003). Consumer and market drivers of the trial probability of new consumer packaged goods. *Journal of Consumer Research* 30(3), 368–384.
- United Nations (2002). *World Population Ageing 1950–2050*. New York: United Nations, Department of Economic and Social Affairs, Population Division.
- Yaari, M. E. (1965). Uncertain lifetime, life insurance, and the theory of the consumer. *Review of Economic Studies* 32(2), 137–150.
- Zajonc, R. B. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology* 9(2, Pt. 2), 1–27.

A Additional figures and tables

Figure 14: Growth and the median age of the population across rich countries




The figure shows the growth of real GDP per capita against the median age of the population across countries, both net of the size and the growth of the working-age population: the 32 OECD members with data since 1970 in panel (a) and the 30 countries of the UN's more developed regions in panel (b). Growth is the average annual rate over 1970–2019 and median age its 1970–2019 average; the controls are the log of the average population aged 15 to 64 and its average annual growth. Each dot is a country, labelled by its ISO code, and the line is the least-squares fit, with its slope and robust standard error in the legend. The regressions are in Table 13. Sources: Penn World Table 11.0; UN World Population Prospects 2024.

Table 13: Growth and the median age of the population across rich countries

	OECD members		UN more developed regions		Western Europe and offshoots	
	(1)	(2)	(3)	(4)	(5)	(6)
Median age (years)	-0.044 (0.031)	-0.145** (0.069)	-0.078* (0.045)	-0.143* (0.075)	-0.133* (0.069)	-0.194* (0.109)
Log working-age population		-0.066 (0.086)		-0.082 (0.113)		-0.213* (0.115)
Growth of the working-age population (pp)		-0.706** (0.359)		-0.674** (0.336)		-0.697* (0.403)
Workforce controls	No	Yes	No	Yes	No	Yes
Countries	32	32	30	30	23	23
R^2	0.07	0.11	0.08	0.21	0.18	0.38


Notes. Least-squares regressions across countries of the average annual growth of real GDP per capita over 1970–2019, in percentage points, on the 1970–2019 average median age of the population. Columns (2), (4) and (6) add the controls of Figure 14. OECD members: current members; UN more developed regions: Europe, Northern America, Australia, New Zealand and Japan; Western Europe as defined by the Maddison Project, with the United States, Canada, Australia and New Zealand. Countries with data since 1970. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Sources: as in Figure 14; Maddison Project Database 2023.

Figure 15: Additional ageing figures



(a) Mean age by data source

(b) Mean age Census projections



(c) Young adults living in the parental home

(a) Mean age, 1980–2025, in the three sources used in the paper. (b) The census series extended by the Census Bureau 2023 national projections: dashed is the middle-immigration series, the band spans the high- and low-immigration series. (c) Share of young adults living in the parental home. Solid lines are the Current Population Survey series; dormitory students count as living at home until 1980 but not later. Dashed lines are the same share in the CEX, restricted to children of the reference person; the CEX treats a dormitory student as a consumer unit

Figure 16: New products by year and brand-age numbers




Table 14: New products and spending by department


Department	Launches	Share of launches (%)	Share of spending (%)	
			Weighted	Unweighted
Dry grocery	31,022	27.9	39.6	39.3
Alcoholic beverages	27,227	24.5	3.8	3.9
Health and beauty care	16,461	14.8	10.5	10.9
General merchandise	13,364	12.0	7.0	7.2
Non-food grocery	10,302	9.3	12.5	12.1
Frozen foods	4,606	4.1	8.7	8.5
Dairy	2,901	2.6	8.9	9.0
Fresh produce	2,673	2.4	3.2	3.4
Deli	1,417	1.3	2.2	2.2
Packaged meat	1,029	0.9	3.6	3.5
Total	111,002	100.0	100.0	100.0

Table 15: Other terms of the regression of Figure 2

	Coefficient	Standard error
<i>Income (omitted: \$50,000–60,000)</i>		
Under \$5,000	−0.0347***	(0.0095)
\$5,000–8,000	−0.0336**	(0.0123)
\$8,000–10,000	−0.0443***	(0.0106)
\$10,000–12,000	−0.0189*	(0.0089)
\$12,000–15,000	−0.0205**	(0.0079)
\$15,000–20,000	−0.0304***	(0.0063)
\$20,000–25,000	−0.0175**	(0.0059)
\$25,000–30,000	−0.0170**	(0.0058)
\$30,000–35,000	−0.0155**	(0.0056)
\$35,000–40,000	−0.0085	(0.0058)
\$40,000–45,000	−0.0074	(0.0056)
\$45,000–50,000	−0.0150**	(0.0055)
\$60,000–70,000	0.0063	(0.0055)
\$70,000–100,000	0.0129**	(0.0047)
\$100,000 and over	0.0440***	(0.0050)
<i>Education of the younger head (omitted: high-school graduate)</i>		
Grade school	0.0197	(0.0169)
Some high school	−0.0000	(0.0059)
Some college	0.0040	(0.0029)
College graduate	0.0141***	(0.0035)
Post-college	0.0134**	(0.0052)
<i>Children and household size (omitted: one member)</i>		
Children under 18	−0.0267***	(0.0043)
Two members	0.0029	(0.0047)
Three members	−0.0088	(0.0056)
Four members	−0.0127*	(0.0065)
Five or more members	−0.0168*	(0.0071)
<i>Years in the panel (omitted: third year or later)</i>		
First year	−0.0040	(0.0034)
Second year	0.0005	(0.0038)
Second year, already in the panel in 2004	−0.0205	(0.0116)
<i>Shopping behaviour</i>		
Log number of shopping trips	0.0143***	(0.0039)
Share of spending bought on deal	−0.0029	(0.0075)
Share of spending saved with coupons	0.3038***	(0.0403)
Price index, %	−0.0020***	(0.0003)
<i>Exposure to new goods</i>		
E_{it}	0.3913***	(0.0271)
Household-years		1,163,731
Households		219,109


Notes. Coefficients of equation (1) estimated on all households with all its terms, as in Figure 2, in percentage points of the share of spending on new products. Shares of spending are fractions between 0 and 1; the price index and the exposure E_{it} are in percent. Years in the panel are counted from 2004, the first year of the data, and households already in the panel in 2004 have their own indicator in their second recorded year. Also included, not shown: the single-year age effects, calendar-year fixed effects, birth-year effects, the employment status of the male and the female head, an indicator for each being 65 or older and their interactions, an indicator for education not stated (2 household-years), the logs of the number of stores and retail chains visited, the share of spending on store brands, the shares of spending by retailer channel, and the logs of the number of distinct products and product modules bought and of total spending. Standard errors clustered by household. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Figure 17: The transition after a one-time jump in age at unchanged familiarity



At date 0 households become as old as in the old era but keep their young-era familiarity with existing products. Since the death hazard does not depend on age, age matters only through experience, the time a household has had to become familiar with products: the young-era distribution of experience converges to the old-era one at the old-era spending hazard (Table 5), and the product vintages are those of the young balanced growth path. The transition comes only from familiarity, and an entrant's sales at launch do not fall on impact. Solid: $\tilde{\rho}$ moves to its old-era value at date 0; dashed: $\tilde{\rho}$ stays at its young-era value. With $\tilde{\rho}$ fixed, growth rises from 2.10% to 2.13% and stays above its initial value for five and a half years: an entrant sells at launch as much as before but expects to sell more once households have become more familiar with its product. With $\tilde{\rho}$ moving, the higher discount outweighs this premium and growth falls at once. *Left*: mean adult age of people and of dollars spent (solid), mean experience of dollars spent (dashed). *Right*: GDP per capita relative to the young balanced growth path, in percent. Grey dashed lines: young balanced growth path.

Figure 18: The transition without foresight



The path of Figure 7 with the mean adult age of expenditure frozen at its 2025 value. Solid: perfect foresight. Dashed: no foresight: at the start of each year innovators observe the current spending hazard $\tilde{\mu}_t$ and expenditure profile ϕ_t , believe that both will stay there forever, and understand that the age distribution and familiarity keep adjusting; the realised path is the first year of the equilibrium path under each year's belief. *Left*: the long-run growth rate innovators expect each year; under perfect foresight, the true one. Without foresight the early-1980s expansion almost disappears: growth peaks at 2.025%, against 2.049%, a small rise left by the ageing that the current hazard alone implies. Without ageing, GDP per capita would be 3.6% higher in 2025 (3.5% under perfect foresight) and 8.1% in 2050 (8.0%).

B Parameters and restrictions

Table 16: Structural parameters

Symbol	Description	Domain
ρ	Subjective discount rate	> 0
μ	Mortality hazard	> 0
ϕ	Rate at which expenditure rises with age across households	$(-\rho, \mu)$
N	Population	> 0
η	Elasticity of substitution across goods (CES)	> 1
g_q	Innovation quality step	> 0
g_f	Familiarity growth rate	> 0
f_0	Baseline familiarity of product–consumer pair	> 0
ν	Entry cost per flow unit rate	> 0
$\bar{\nu}$	Composite entry cost per unit rate	$= \eta\nu/N$
$\tilde{\mu}$	Spending hazard	$= \mu - \phi$
$\tilde{\rho}$	Gap between the return and growth, $r - g$	$= \rho + \phi$

Note: the scale of the economy N and the baseline familiarity level f_0 affect the equilibrium only through the composite $\bar{\nu} = \eta\nu/N$; the balanced growth path and its comparative statics therefore depend on the seven parameters $\{\rho, \mu, \phi, \eta, g_q, g_f, \bar{\nu}\}$; the rates of creative destruction and growth depend on ρ, μ and ϕ only through $\tilde{\mu}$ and $\tilde{\rho}$.

Three parameter restrictions.

- **Expenditure profile:** $-\rho < \phi < \mu$. The upper bound makes the income profile (6) integrable; the lower bound keeps the gap between the return and growth, $r - g = \rho + \phi$, positive.
- **Familiarity boundedness:** $\min\{g_f, (1 + g_q)\tau\} < \tilde{\mu}$. Ensures $\mathcal{F}(\infty) < \infty$, but is not satisfied in the calibrations in the quantitative part of the paper. When $g_f < (1 + g_q)\tau$, this reduces to $g_f < \tilde{\mu}$.
- **Finite entry and incumbent value:** $\min\{g_f, (1 + g_q)\tau\} < \rho + \mu + (1 + g_q)\tau$. For the problem to be well defined we need the value of a firm to be finite. The condition comes from the Laplace transform $\widehat{\mathcal{F}}(\tilde{\rho} + (1 + g_q)\tau)$, and is in fact no restriction at all: its left-hand side is at most $(1 + g_q)\tau$, so it holds with slack $\rho + \mu > 0$ (note $\tilde{\rho} + \tilde{\mu} = \rho + \mu$) at every admissible parameter configuration, independently of the two conditions above. Proposition 3 sharpens this into the bound (35): for every $\tau > 0$ the entry value is finite and below $1/(\tilde{\rho} + (1 + g_q)\tau)$.

C Proofs

The proofs of Lemmas 1–2 and Propositions 1–4 set $\phi = 0$, so that $\tilde{\mu} = \mu$ and $\tilde{\rho} = \rho$. Since the expenditure profile enters these results only through $\tilde{\mu}$ and $\tilde{\rho}$, they hold for any $\phi \in (-\rho, \mu)$ with $(\tilde{\mu}, \tilde{\rho})$ in place of (μ, ρ) . Throughout, $\sigma(\tau) \equiv \rho + (1 + g_q)\tau$.

Proof of Lemma 1. Take any continuous path τ_t . Over an interval Δt , a fraction $\tau_t \Delta t$ of lines innovates, multiplying q_j by $1 + g_q$, and lines that innovate more than once have measure $o(\Delta t)$. Hence $\bar{Q}_{t+\Delta t}^{\eta-1} = (1 + g_q \tau_t \Delta t) \bar{Q}_t^{\eta-1} + o(\Delta t)$, and letting $\Delta t \rightarrow 0$ gives $\dot{\bar{Q}}_t / \bar{Q}_t = g_q \tau_t / (\eta - 1)$. $\square$

Proof of Lemma 2. A line of vintage age a^p at t last innovated at $t - a^p$. Creative destruction strikes lines independently of their quality, so its quality before that innovation is a draw from the cross-section at date $t - a^p$, with mean $\bar{Q}_{t-a^p}^{\eta-1}$, and survival to t involves no selection. The innovation multiplies quality by $1 + g_q$, so

$$Q_t(a^p)^{\eta-1} = (1 + g_q) \bar{Q}_{t-a^p}^{\eta-1} = (1 + g_q) e^{-(\eta-1)g_q a^p} \bar{Q}_t^{\eta-1},$$

where the last equality holds on the balanced path. $\square$

The remaining proofs use the household average familiarity $\mathcal{F}^h$ of (24). Splitting the average at the household's age gives, for all parameter values,

$$\frac{\mathcal{F}^h(a^h)}{f_0} = 1 + g_f \int_0^{a^h} e^{-[(1+g_q)\tau - g_f]v} dv. \quad (65)$$

Hence $\mathcal{F}^h(0) = f_0$, and $\mathcal{F}^h$ increases with age, with $(\mathcal{F}^h)'(a^h) = g_f f_0 e^{-[(1+g_q)\tau - g_f]a^h}$. The product index is

$$\mathcal{F}(a^p) = \int_0^\infty \frac{f_0 e^{g_f \min\{a^h, a^p\}}}{\mathcal{F}^h(a^h)} \mu e^{-\mu a^h} da^h. \quad (66)$$

Averaging it over product ages with the quality-weighted density $\omega(a^p) \equiv (1 + g_q)\tau e^{-(1+g_q)\tau a^p}$ and exchanging the integrals gives the normalisation

$$\int_0^\infty \mathcal{F}(a^p) \omega(a^p) da^p = \int_0^\infty \frac{\mathcal{F}^h(a^h)}{\mathcal{F}^h(a^h)} \mu e^{-\mu a^h} da^h = 1. \quad (67)$$

Proof of Proposition 1. The proposition holds for all parameter values. By (65), $f_0 \leq \mathcal{F}^h(a^h) < \infty$ at every age, and $\mathcal{F}^h$ converges to $\mathcal{F}^h(\infty) = f_0(1 + g_q)\tau / [(1 + g_q)\tau - g_f]$ if $g_f < (1 + g_q)\tau$ and diverges otherwise. Hence (66) gives $\mathcal{F}(a^p) \leq e^{g_f a^p} \mathcal{F}(0) \leq e^{g_f a^p}$.

(i) Splitting (66) at $a^h = a^p$ and differentiating, the boundary terms cancel, and only households older than the product contribute to its familiarity growth:

$$\mathcal{F}'(a^p) = g_f \mu f_0 e^{g_f a^p} \int_{a^p}^{\infty} \frac{e^{-\mu a^h}}{\mathcal{F}^h(a^h)} da^h > 0.$$

At introduction all households contribute, so $\mathcal{F}'(0) = g_f \mathcal{F}(0)$. At positive ages, households younger than the product contribute to $\mathcal{F}$ but not to $\mathcal{F}'$, so $0 < \mathcal{F}'/\mathcal{F} < g_f$.

(ii) Since $\mathcal{F}$ is strictly increasing with ω -average one, $\mathcal{F}(0) < 1 < \mathcal{F}(\infty)$, and by monotone convergence $\mathcal{F}(\infty) = \mu f_0 \int_0^{\infty} e^{(g_f - \mu)a^h} / \mathcal{F}^h(a^h) da^h$. If $g_f < (1 + g_q)\tau$, $\mathcal{F}^h$ lies between f_0 and $\mathcal{F}^h(\infty)$, so the integral is finite exactly when $g_f < \mu$. If $g_f > (1 + g_q)\tau$, (65) bounds $\mathcal{F}^h/f_0$ above and below by constants times $e^{[g_f - (1 + g_q)\tau]a^h}$, so the integral is finite exactly when $(1 + g_q)\tau < \mu$. At equality, $\mathcal{F}^h/f_0 = 1 + g_f a^h$, and the integral is finite exactly when $g_f < \mu$. Hence $\mathcal{F}(\infty) < \infty$ if and only if $\min\{g_f, (1 + g_q)\tau\} < \mu$.

(iii) Only the age density depends on μ . Differentiating (66),

$$\partial_{\mu} \mathcal{F}(a^p) = f_0 \int_0^{\infty} \frac{e^{g_f \min\{a^h, a^p\}}}{\mathcal{F}^h(a^h)} (1 - \mu a^h) e^{-\mu a^h} da^h, \quad (68)$$

and, at $a^p = 0$, integrating by parts,

$$\partial_{\mu} \mathcal{F}(0) = f_0 \int_0^{\infty} \frac{a^h (\mathcal{F}^h)'(a^h)}{[\mathcal{F}^h(a^h)]^2} e^{-\mu a^h} da^h > 0: \quad (69)$$

a younger population gives more weight to households with less familiarity to compete against. Differentiating (68) with respect to product age,

$$\partial_{a^p} \partial_{\mu} \mathcal{F}(a^p) = g_f f_0 e^{g_f a^p} \int_{a^p}^{\infty} \frac{(1 - \mu a^h) e^{-\mu a^h}}{\mathcal{F}^h(a^h)} da^h.$$

The tail integral is positive at zero by (69), decreases up to $a^p = 1/\mu$, where it is negative, and then rises towards zero. So $\partial_{\mu} \mathcal{F}$ first increases and then decreases. Since (67) holds for every μ and ω does not depend on μ , $\partial_{\mu} \mathcal{F}$ averages to zero over product ages, so it crosses zero exactly once, from positive to negative.

(iv) By (65),

$$\partial_{\tau} \mathcal{F}^h(a^h) = -f_0 g_f (1 + g_q) \int_0^{a^h} v e^{-[(1 + g_q)\tau - g_f]v} dv < 0 \quad \text{for } a^h > 0.$$

In (66), τ enters only through $\mathcal{F}^h$, so

$$\partial_\tau \mathcal{F}(a^p) = -\mu f_0 \int_0^\infty \frac{e^{g_f \min\{a^h, a^p\}} \partial_\tau \mathcal{F}^h(a^h)}{[\mathcal{F}^h(a^h)]^2} e^{-\mu a^h} da^h > 0.$$

Differentiation under the integral is legitimate, since $|\partial_\tau \mathcal{F}^h(a^h)| \leq (1 + g_q) a^h \mathcal{F}^h(a^h)$. $\square$

Proof of Proposition 2. At fixed quality, profits (27) grow at $\mathcal{F}'/\mathcal{F} + (2 - \eta)g$. Substituting $V = \kappa \pi_j$ into the HJB and dividing by profits gives $[r + \tau - (2 - \eta)g]\kappa - \kappa' - \kappa \mathcal{F}'/\mathcal{F} = 1$, which is (29) since $r = \rho + g$ and $(\eta - 1)g = g_q \tau$. This is a linear equation for $\kappa \mathcal{F}$, with general solution

$$\kappa(a^p) \mathcal{F}(a^p) = C e^{\sigma(\tau) a^p} + \int_0^\infty e^{-\sigma(\tau) u} \mathcal{F}(a^p + u) du.$$

The integral is finite, since (67) and $\sigma(\tau) > (1 + g_q)\tau$ imply $\widehat{\mathcal{F}}(\sigma(\tau)) \leq 1/[(1 + g_q)\tau]$. Transversality holds for the integral solution and fails for every $C \neq 0$, which gives (30).

(i) By Proposition 1(i), $\mathcal{F}(a^p + u)/\mathcal{F}(a^p) \leq e^{g_f u}$. If $g_f < (1 + g_q)\tau$, the integrand of (30) and its age derivative are therefore bounded by $e^{-(\sigma(\tau) - g_f)u}$ and $2g_f e^{-(\sigma(\tau) - g_f)u}$, and differentiating under the integral gives (31).

(ii) The ratio $\mathcal{F}'/\mathcal{F}$ equals g_f at introduction and is smaller at every positive age, so the bracket in (31) is negative and $\kappa'(0) < 0$. $\square$

Proof of Proposition 3. Exchange the integrals over product and household ages in (66), which is legitimate since the integrand is positive. For each household age, integration by parts gives

$$\sigma(\tau) \int_0^\infty e^{-\sigma(\tau) u} e^{g_f \min\{a^h, u\}} du = 1 + g_f \int_0^{a^h} e^{-(\sigma(\tau) - g_f)v} dv,$$

so that, by (65), the entry value is

$$\widehat{\mathcal{F}}(\sigma(\tau); \mu, \tau) = \frac{1}{\sigma(\tau)} \int_0^\infty \frac{1 + g_f \int_0^{a^h} e^{-[\rho + (1 + g_q)\tau - g_f]v} dv}{1 + g_f \int_0^{a^h} e^{-[(1 + g_q)\tau - g_f]v} dv} \mu e^{-\mu a^h} da^h. \quad (70)$$

The ratio compares the same household's familiarity with and without the additional discount ρ . It lies in $(0, 1]$, equals one if $g_f = 0$, and is below one at positive ages if $g_f > 0$. This proves (35), finiteness for every $\tau > 0$ and the zero limit as $\tau \rightarrow \infty$. The ratio is continuous in τ and bounded by one, so by dominated convergence the entry value is continuous, including as $\tau \downarrow 0$. In that limit the denominator becomes $e^{g_f a^h}$, and integrating the limiting ratio against $\mu e^{-\mu a^h}$ and dividing by ρ gives (36).

The entry value is bounded by $1/\rho$; let $\bar{V}$ be its supremum. If $\bar{v}/(1 + g_q) < \bar{V}$, the entry value

exceeds the cost at some rate and falls below it at large rates, so (33) has a solution. At every solution, $\bar{v}\tau = (1+g_q)\tau\widehat{\mathcal{F}}(\sigma(\tau);\mu,\tau) \leq (1+g_q)\tau/\sigma(\tau) < 1 < \eta$, so feasibility holds. If $\bar{v}/(1+g_q) > \bar{V}$, no positive-entry BGP exists. $\square$

Proof of Proposition 4. At $g_f = 0$ the ratio in (70) is one, so the entry value is $1/\sigma(\tau)$, and its derivative $-(1+g_q)/\sigma(\tau)^2$ is bounded away from zero for $0 \leq \tau \leq \eta/\bar{v}$. For $0 \leq g_f \leq \mu/2$ and τ in this interval, the ratio and its τ -derivative are jointly continuous and bounded by a constant times $(1+(a^h)^2)e^{\mu a^h/2}$, which is integrable against $\mu e^{-\mu a^h}$. By dominated convergence, the derivative of the entry value is uniformly continuous on this compact set, so it stays negative for g_f in some interval $[0, \bar{g}_f]$. The entry value then strictly decreases on the feasible interval and equals the entry cost at most once. $\square$

Proof of Lemma 3. Making the quality weights explicit, $\mathcal{F}_t^h(a^h) = \int_0^1 q_{jt} f_{ijt} dj / \bar{Q}_t^{\eta-1}$. Conditioning on product age gives (24), and $\mathcal{F}_t^h(0) = f_0$ at every date. Along a household's life, familiarity grows at g_f on lines that do not innovate, and innovation at rate τ_t replaces $q_j f_{ij}$ with $(1+g_q)q_j f_0$. Since $\bar{Q}_t^{\eta-1}$ grows at $g_q\tau_t$ (Lemma 1),

$$(\partial_t + \partial_{a^h})\mathcal{F}_t^h(a^h) = [g_f - (1+g_q)\tau_t]\mathcal{F}_t^h(a^h) + (1+g_q)\tau_t f_0.$$

Dividing by $\mathcal{F}_t^h$ gives the law of motion of a household's average familiarity over her life, along any equilibrium path:

$$(\partial_t + \partial_{a^h})\log \mathcal{F}_t^h(a^h) = g_f - (1+g_q)\tau_t \left(1 - \frac{f_0}{\mathcal{F}_t^h(a^h)}\right). \quad (71)$$

Let $\Gamma_t \equiv \int_0^\infty \psi_t(a^h) \log(\mathcal{F}_t^h(a^h)/f_0) da^h$ be its population average, where the age density obeys $(\partial_t + \partial_{a^h})\psi_t = -\mu_t\psi_t$, with $\psi_t(0) = \mu_t$. Since $0 \leq \log(\mathcal{F}_t^h/f_0) \leq g_f a^h$ and, starting from a BGP, the age density has an exponential tail, Γ_t is finite. Differentiating it and integrating by parts, with no birth term since $\mathcal{F}_t^h(0) = f_0$,

$$\dot{\Gamma}_t = \int_0^\infty \psi_t(a^h)(\partial_t + \partial_{a^h})\log \frac{\mathcal{F}_t^h(a^h)}{f_0} da^h - \mu_t \Gamma_t.$$

Substituting (71), the population average of $f_0/\mathcal{F}_t^h$ is $\mathcal{F}_{\mu,t}(0)$, the familiarity index of a new product with every household weighted equally. Hence

$$\dot{\Gamma}_t = g_f - (1+g_q)\tau_t(1 - \mathcal{F}_{\mu,t}(0)) - \mu_t \Gamma_t.$$

On a BGP, $\dot{\Gamma} = 0$ and $\psi_t(a^h) = \mu e^{-\mu a^h}$, which gives (52). $\square$

Proof of Proposition 7. Fix τ and $\mu_1 < \mu_2$, and write $\mathbb{E}_{\mu_i}$ for expectations under $\text{Exp}(\mu_i)$. The familiarity term of (51) depends on μ only through the age distribution, and, integrating by parts with $\log(\mathcal{F}^h(0)/f_0) = 0$,

$$\mathbb{E}_{\mu_1} \left[\log \frac{\mathcal{F}^h(a^h)}{f_0} \right] - \mathbb{E}_{\mu_2} \left[\log \frac{\mathcal{F}^h(a^h)}{f_0} \right] = \int_0^\infty \frac{(\mathcal{F}^h)'(a^h)}{\mathcal{F}^h(a^h)} (e^{-\mu_1 a^h} - e^{-\mu_2 a^h}) da^h > 0,$$

since $(\mathcal{F}^h)' > 0$ and the survival gap is positive on $(0, \infty)$. At $\phi = 0$ the life-cycle term is zero and the familiarity term is the only one that depends on μ , so $\mathcal{W}(\mu_1, \tau) > \mathcal{W}(\mu_2, \tau)$. Since $\mathcal{W}(\mu, \cdot)$ is continuous on $[0, \eta/\bar{v})$ and diverges to $-\infty$ at the right endpoint, its maximum is attained, say at $\tau^{\mathcal{W}}(\mu_2)$, and

$$\mathcal{W}^*(\mu_2) = \mathcal{W}(\mu_2, \tau^{\mathcal{W}}(\mu_2)) < \mathcal{W}(\mu_1, \tau^{\mathcal{W}}(\mu_2)) \leq \mathcal{W}^*(\mu_1).$$

□

Proof of Proposition 8. Every good costs one unit of labour, so research is sourced by minimising $\int_0^1 x_j dj$ subject to

$$\left(\int_0^1 w_j^{1/\eta} x_j^{(\eta-1)/\eta} dj \right)^{\eta/(\eta-1)} \geq v\tau \bar{Q}, \quad \int_0^1 w_j dj = \bar{Q}^{\eta-1}.$$

The CES first-order conditions give $x_j \propto w_j$, hence $x_j = v\tau w_j / \bar{Q}^{\eta-1}$ and $L^R = v\tau$. This holds for the planner's weights $w_j = q_j$ and for the market's, $q_j \mathcal{F}(a_j^p)$, which satisfy the normalisation by (67). Research thus leaves $N - v\tau$ goods for consumption. With multiplier λ on this constraint, the first-order condition for each household and product gives $c_{ij} = q_j f_{ij} C_i^{1-\eta} \lambda^{-\eta}$. Substituting into the aggregator and using $\int_0^1 q_j f_{ij} dj = \bar{Q}_t^{\eta-1} \mathcal{F}_t^h(a_i^h)$ gives $C_i = \lambda^{-1} \bar{Q}_t [\mathcal{F}_t^h(a_i^h)]^{1/(\eta-1)}$ and $\int_0^1 c_{ij} dj = \lambda^{-1}$: every household gets the same quantity of goods. Market clearing gives $1/\lambda = 1 - \bar{v}\tau/\eta$, which proves (56).

When $\phi = 0$, the market delivers the same allocation. All goods have the same price, so uniform markups do not distort relative prices; every household buys $E/p = 1 - \bar{v}\tau/\eta$ goods by (34) and allocates them in proportion to $q_j f_{ij}$; and research is sourced as above. Taking logs of (56) and averaging across households gives (49). □

Proof of Proposition 9. With the static allocation (56) and Lemma 1, welfare (47) becomes

$$\frac{\mathcal{W}}{N} = \int_0^\infty e^{-\rho t} \left[\log \left(1 - \frac{\bar{v}\tau}{\eta} \right) + \log \bar{Q}_0 + \frac{g_q}{\eta-1} \int_0^t \tau_s ds + \frac{\log f_0 + \Gamma_t}{\eta-1} \right] dt, \quad (72)$$

where the initial quality and familiarity distributions are held fixed.

(i) Add a uniform innovation mass εds at date s to the constant path, starting from its stationary cross-section. At $s + u$, the new vintage has age u and tilted mass $(1 + g_q)\varepsilon ds e^{-(1+g_q)\tau u}$, while vintages older than u lose the fraction $(1 + g_q)\varepsilon ds$ of their tilted mass (Appendix D). Since Γ_t averages $\log(\mathcal{F}_t^h/f_0)$ over households, its first-order change averages $d\mathcal{F}_t^h/\mathcal{F}_t^h$. By (66), the new vintage therefore contributes $\mathcal{F}(u)$ and the displaced vintages their ω -weighted index:

$$d\Gamma_{s+u} = \varepsilon ds(1 + g_q) \left[e^{-(1+g_q)\tau u} \mathcal{F}(u) - \int_u^\infty \mathcal{F}(a^p) \omega(a^p) da^p \right].$$

Discounted at ρ , the first term gives $\widehat{\mathcal{F}}(\sigma(\tau))$. By (67), the second gives $[1 - (1 + g_q)\tau \widehat{\mathcal{F}}(\sigma(\tau))]/\rho$. Hence

$$\int_0^\infty e^{-\rho u} d\Gamma_{s+u} du = \varepsilon ds(1 + g_q) \frac{\sigma(\tau) \widehat{\mathcal{F}}(\sigma(\tau)) - 1}{\rho} < 0, \quad (73)$$

where the sign follows from (35) for $g_f > 0$. Per unit of $e^{-\rho s} N \varepsilon ds$, the impulse changes the resource term by $-\bar{v}/(\eta - \bar{v}\tau)$ at s , raises the growth term by $g_q/[(\eta - 1)\rho]$ in present value, and changes the familiarity term by $1/(\eta - 1)$ times (73). Setting the sum to zero and multiplying by $(\eta - 1)\rho$ gives (57).

(ii) At $g_f = 0$, $\Gamma_t \equiv 0$, so (72) is the sum of a strictly concave and a linear functional of the path. The pointwise first-order condition $\bar{v}/(\eta - \bar{v}\tau_t) = g_q/[(\eta - 1)\rho]$ is then sufficient and gives (58), truncated at zero.

(iii) By (36), as $\tau \downarrow 0$ the left-hand side of (57) tends to $(1 + g_q)[1 - \rho g_f/((\mu + g_f)(\rho + \mu))]$ $- 1$, and the right-hand side to $(\eta - 1)\rho \bar{v}/\eta$. Under (59), the social value of an innovation therefore falls short of its cost on a neighbourhood $(0, \bar{\tau})$, which contains no interior stationary first best. For stationary welfare, $\mathcal{F}(0) \rightarrow \mu/(\mu + g_f)$ as $\tau \downarrow 0$, and (53) gives

$$\partial_\tau \mathcal{W}|_{\tau=0} = -\frac{\bar{v}}{\eta} + \frac{g_q}{(\eta - 1)\rho} - \frac{(1 + g_q)g_f}{(\eta - 1)\mu(\mu + g_f)}.$$

Its familiarity loss exceeds that of the impulse derivative, whose denominator is $(\eta - 1)(\rho + \mu)(\mu + g_f)$, so (59) also makes it negative near zero. Neither argument establishes optimality against arbitrary changes in the innovation path. $\square$

Proof of Proposition 10. Differentiating (57) with respect to μ , with Φ as in Proposition 10, gives

$$\left[\frac{(\eta - 1)\rho \bar{v}^2}{(\eta - \bar{v}\tau)^2} - \Phi_\tau \right] \frac{d\tau^{\text{FB}}}{d\mu} = (1 + g_q)\sigma(\tau) \partial_\mu \widehat{\mathcal{F}}.$$

The right-hand side is positive under $g_f < (1 + g_q)\tau$, by the argument in the proof of Proposition 5, which concerns only the entry-value function. On the left, (35) gives $\sigma(\tau) \widehat{\mathcal{F}} < 1$ when $g_f > 0$, so at any interior first best (57) implies $\tau^{\text{FB}} < \eta/\bar{v} - (\eta - 1)\rho/g_q$. Restrict τ to the interval

between zero and this bound, and μ to $[\mu_1, \mu_2]$. As in the proof of Proposition 4, (70) makes Φ_τ uniformly continuous over these intervals and $0 \leq g_f \leq \mu_1/2$. At $g_f = 0$, $\Phi \equiv 1 + g_q$ and $\Phi_\tau = 0$, while the resource slope is at least $(\eta - 1)\rho\bar{v}^2/\eta^2 > 0$. Hence there is a threshold $\check{g}_f > 0$ below which the coefficient on the left is positive, and the implicit function theorem gives (63). $\square$

D Computing the transition

This appendix describes how we compute the equilibrium transition between balanced growth paths: an unanticipated, permanent fall in the death hazard from μ^{young} to μ^{old} at $t = 0$, with natality adjusting one-for-one, starting on the young-era balanced growth path. Off the balanced path the household age and vintage distributions become state variables, both with closed forms: the age density splices the new-hazard exponential below the shock-cohort age onto the thinned pre-shock tail above it, and the tilted density of Lemma 1 generalises to $\omega_u(a^p) = (1 + g_q)\tau_{u-a^p}e^{-(1+g_q)[T(u)-T(u-a^p)]}$, with $T(t) \equiv \int_0^t \tau_s ds$ and τ_s equal to the young-era rate for $s < 0$. The date- u familiarity index $\mathcal{F}_u$ is then (22) evaluated at the transiting densities (familiarity is still (3)), and expenditure stays common across cohorts (the hazard is age-independent), proportional to output net of entry costs. The equilibrium is a path $\{\tau_t\}_{t \geq 0}$ satisfying free entry at every date:

$$Z_t \equiv (1 + g_q) \int_t^\infty e^{-\rho(u-t)-(1+g_q)[T(u)-T(t)]} \frac{1 - \bar{v}\tau_t/\eta}{1 - \bar{v}\tau_u/\eta} \mathcal{F}_u(u-t) du = \bar{v} \quad \text{for all } t \geq 0, \quad (74)$$

which collapses to (33) on a balanced path. We solve the fixed point as follows:

- (i) Discretise time and both age dimensions on a common uniform grid of step $\Delta = 0.5$ years over a 250-year horizon, truncating the entrant's life in (74) at 80 years with an analytic old-era continuation beyond it;
- (ii) Compute the discretised anchors: the entry value on the flat young-era path defines the numerical entry cost $\bar{v}_\Delta$, and a secant iteration on the flat old-era map delivers the discretised root τ_Δ^{old} ; using these in place of their exact counterparts makes the no-shock path exactly flat, so the discretisation bias cancels from the fixed point, and the final 80 years are pinned at τ_Δ^{old} (the terminal condition);
- (iii) Given a candidate path, rebuild the densities and $\mathcal{F}_u$ and evaluate Z_t at every unpinned date. The densities jump at the shock-cohort boundary, where the grid node takes the midpoint of the one-sided limits; without this rule the quadrature adds spurious mass of order Δ that distorts the impact response;

- (iv) Solve $Z_t = \bar{v}_\Delta$ by Newton’s method in sequence space (Auclert et al., 2021): the Jacobian $\partial Z_t / \partial \tau_j$, zero beyond the entrant horizon, is built once by forward differences around the flat old-era path, factorised, and reused at every backtracked Newton step until the maximum relative residual falls below 10^{-10} .

In addition, we perform a battery of tests to check the algorithm performs correctly under known special cases: a constant- μ placebo returns a flat path; the long-run limit reproduces the closed-form old-era solution of (33); the path version of the normalisation (23) holds; the path is flat as $g_f \rightarrow 0$ (μ -neutrality); we also check that halving Δ leaves the solution unchanged.

E New products from new firms

In Section 2, I define a product as a brand-module pair and select new products among brand-module debuts. This appendix compares my launches with a definition based on the entry of new firms. It relies on barcodes. Every barcode starts with a company prefix, which GS1 assigns to the firm that owns the barcode. A barcode whose prefix has never been seen before belongs to a firm that is new to the market.

NielsenIQ does not record where the company prefix ends, and prefixes differ in length. I take the length from the GS1 Company Prefix length table, which gives the prefix length for each range of barcodes.⁵² Most barcodes in the panel (78%) carry six-digit prefixes, the original length of UPC company prefixes. About 11% carry prefixes of seven to eleven digits, which leave room for fewer products and are typical of smaller firms.

I date the entry of a prefix to the first year in which a panellist buys any barcode carrying it, and I treat 2004 as a burn-in year. I call new-firm goods the barcodes that debut in the entry year of their prefix. Established firms also buy new prefixes, for example when they run out of product numbers. I therefore drop a new prefix if any of the brands it sells in its entry year had already been bought under another prefix. As in Section 2, I exclude private labels, pseudo-brands and manufacturer names without a brand. NielsenIQ replaces the barcodes of private labels with codes that carry no information on the producer.

Panel A of Table 17 reports the counts. Between 2005 and 2024, 2,008,463 branded barcodes debut in the panel. Of these, 128,551 debut in the entry year of a new prefix. After dropping the prefixes that also sell an older brand, I am left with 69,236 new-firm goods, sold by 38,288 new firms. This is about 3,500 goods a year, 3.4% of barcode debuts and 1.3% of the spending on them in their first year. The share is stable over time: it moves between 2.3% and 4.2% of debuts in every year except 2009 (6.2%).

⁵²GS1, “GCP length table”, version of 17 February 2026, from GS1’s public repository <https://github.com/gs1/TDT>. The table gives the length of the prefixes, not the names of the firms that own them.

Two caveats apply. First, entry is dated by the first purchase in the panel, not by the founding of the firm, so a small firm can enter the panel long after it started selling. Second, the rule misses established firms that launch a brand under a new name with a newly bought prefix, such as Abbott’s BinaxNOW Covid test in 2021 or the Fitbit Charge 3 in 2018. These counts are therefore an upper bound on the goods of new firms.

Table 17: New-firm goods in the Homescan panel

<i>Panel A. From barcode debuts to new-firm goods, 2005–2024</i>				
	Barcodes	GS1 prefixes		
Branded barcode debuts	2,008,463			
in the entry year of a new GS1 prefix	128,551	59,381		
less prefixes that also sell an older brand	69,236	38,288		
<i>Panel B. New-firm goods by step of Table 2</i>				
	Barcodes	Barcodes (%)	Spending (%)	Brand-modules
In the launch sample	59,408	85.8	83.7	39,533
Module's own first year	2,865	4.1	3.2	309
Recently relabelled barcodes	214	0.3	0.3	107
Name extensions of an older brand	6,749	9.7	12.8	4,445
Total	69,236	100.0	100.0	44,394

Notes. Branded barcode debuts are the barcodes first bought by a panellist between 2005 and 2024, excluding private labels, pseudo-brands, manufacturer names without a brand and barcodes without a GS1 prefix. A GS1 prefix enters in the first year in which a panellist buys any barcode carrying it; 2004 is a burn-in year. Prefix lengths are from the GS1 GCP length table. New-firm goods are the barcodes that debut in the entry year of their prefix, when the prefix sells no brand bought earlier under another prefix. Panel B classifies each new-firm good by the brand-module it sells in its first year and by the step of Table 2 at which that brand-module leaves the launch sample. Spending is first-year spending with projection weights.

Panel B of Table 17 shows that the new-firm goods are almost all in my launch sample. 85.8% of them, with 83.7% of their spending, are launch-year barcodes of the 111,002 launches. Most of the remainder (9.7%) are excluded as name extensions of an older brand: they carry names, such as Fitbit Charge 3 or Clorox Pool & Spa, that continue an established brand under a new prefix. Another 4.1% are sold in the first year of a newly introduced module, and only 0.3% are on relabelled barcodes. No new-firm good belongs to a brand already selling in another module, since the rule above excludes them.

Table 18 takes the opposite direction and classifies the launch-year barcodes of my 111,002 launches. A third of the launches (35.1%) come entirely from new firms, but they account for only 22.2% of launch-year spending. Another 26.8% come from young firms, whose prefix entered the panel after 2004 but before the launch. A further third (33.8%) are new brands of firms

already selling in 2004, and these carry more than half (55.7%) of the spending on launches in their first year. The remaining launches are new brands that a firm selling an older brand launches under a new prefix (1.7%), launches on barcodes that sold under another brand or module more than two years earlier (1.0%), launches without a GS1 prefix (0.7%) and launches whose barcodes fall in more than one class (0.9%).

Table 18: Launches by the GS1 status of their barcodes

GS1 status of the barcode	Launch-year barcodes		Launches	
	Barcodes (%)	Spending (%)	Number	Share (%)
New firm: in the entry year of its prefix	35.4	22.2	38,933	35.1
New prefix that also sells an older brand	2.4	2.7	1,925	1.7
Young firm: prefix entered after 2004, before the launch	25.8	18.7	29,720	26.8
Incumbent firm: prefix on sale in 2004	35.2	55.7	37,495	33.8
Barcode sold before the launch year	0.8	0.4	1,100	1.0
No GS1 prefix	0.6	0.3	796	0.7
Barcodes in more than one class	–	–	1,033	0.9
Total	100.0	100.0	111,002	100.0

Notes. Launch-year barcodes are the 167,987 barcodes bought in the launch year of the 111,002 launches. New firm: the barcode debuts in the entry year of its prefix and the prefix sells no older brand (the new-firm goods of Table 17). New prefix that also sells an older brand: the prefixes dropped in Panel A of Table 17. Young firm: the prefix entered between 2005 and the year before the launch. Incumbent firm: the prefix was on sale in 2004. A launch is assigned to a class when all its launch-year barcodes fall in it. Spending is launch-year spending with projection weights.

A definition based on new firms would therefore keep about a third of my launches and a fifth of their spending. It would miss the new brands of established firms, which account for most of the spending on new products. Conversely, almost all the products that new firms bring to the market are already among my launches.

F New product adoption in the CEX

The HMS evidence is rich but comes from one specific kind of consumption: packaged goods sold in supermarkets. To see whether the age gradient extends beyond this restricted sample, I turn to the CEX. The CEX is a widely used consumption survey tracking spending in all product categories, including consumption goods, services, housing and health. It includes the age of all members of the household, whose average tracks very closely the average age in the population, and identifies a reference person in the household whose age I use in the analysis. The reference person is the first person the respondents name among those who own or rent the home. The dataset also includes additional household characteristics such as income and education. Compared to the HMS, it has the advantage of covering all spending categories, but it does not record purchases of individual goods. Spending is instead grouped into categories that coincide with UCC codes, 847 in total over the life of the survey. Additionally, it is only a short panel: each household is interviewed up to four times, once a quarter, about its spending in the previous three months, so it is followed for at most one year.

While individual purchases are not recorded in the CEX, new categories are added over time, and I use them to study how adoption correlates with age. I focus on three specific categories: cellular phones, home internet and video streaming. Since these categories require a subscription, this choice minimises the risk of households forgetting a purchase or having adopted in the past without reporting a purchase in the current year. I define the penetration of a category in a year as the share of consumer units, weighted by their survey weights, that report positive spending on it in any month of that year in which they are observed. I identify each category's base year as the first year in which penetration reaches 1% and is still above it the following year. Cellular telephone service (UCC 270102) first appears in the survey in 1993 and reaches its base year in 1994, when 3.7% of consumer units report it; ten years later the share is 48.6%. Home internet service (UCC 690114) appears in 1995, has its base year in 1996 at 6.3% and reaches 53.8% ten years on. Video streaming (UCCs 310240 and 310243) appears in 2005, touches 1.1% in 2006 but falls back to 0.8% in 2007, and only crosses 1% durably in 2008, at 1.1%; a decade later it stands at 30.0%.⁵³

For each category g , I estimate a separate linear probability model on the consumer units observed in its base year t_g :

$$y_{ig} = \alpha_g + \beta_g \frac{\text{age}_{it_g}}{10} + \sum_{m=2}^{12} \lambda_{gm} \mathbf{1}\{M_{ig} = m\} + \gamma'_g Z_{it_g} + \varepsilon_{ig}. \quad (75)$$

⁵³Each interview asks about spending in the previous three months, so the interviews of the first quarter of a year cover the last quarter of the previous one. A category that enters the questionnaire at the beginning of a year is thus first recorded in the last quarter of the year before. The penetration rule naturally excludes years in which a category was recorded only for a few months.

The outcome y_{ig} equals 100 for adopters and zero otherwise. I count a household as an adopter if it reports positive spending in any observed month of the year. Thus, β_g measures the difference in adoption associated with ten additional years of age, in percentage points. The variable M_{ig} counts the months of the year in which the household is observed in the survey, and its indicators account for the fact that a household that is interviewed multiple times has more opportunities to report a purchase. I first include only age and these indicators, then add controls in Z_{itg} : income-quintile indicators, household size (one to five or more), the number of children under 18 (zero to three or more), education and home ownership. Income quintiles are computed with survey weights within each base year's sample. Age is that of the CEX reference person, and demographics come from the household's latest interview for that year; the sample includes reference persons aged 18 or over. I estimate equation (75) by weighted least squares, using each consumer unit's mean survey weight over its observed quarters. Since each regression uses a single year's cross-section, it does not separately identify age and birth-cohort effects. As in the HMS, I express the ten-year age association relative to the weighted average adoption rate to compare magnitudes across services.

Table 19 reports estimation results. Older households are less likely to use all three services. Without household controls, ten additional years of age are associated with adoption probabilities lower by 11.1%, 12.5% and 21.6% of the average for cellular phones, home internet and video streaming, respectively. Adding controls reduces the differences, but the age gradient remains sizeable: the corresponding declines are 10.5%, 8.1% and 11.5%. In percentage points, these are reductions of 0.38, 0.51 and 0.13. The first two are statistically significant at the 1% level, the third at the 5% level. Differences in resources and household characteristics therefore account for part of the age gradient, but older households remain less likely to use these services after controlling for them.

These results complement the evidence from the HMS. The negative age gradient appears both in purchases of new packaged goods and in the use of new service categories early in their diffusion, and magnitudes are comparable in the two samples: with controls, ten additional years of age reduce adoption by around 10% of the average rate in both datasets. The three cases do not represent the whole consumption basket, but they show that the association between younger age and adoption extends beyond the goods typically purchased at the supermarket.

Table 19: Age and launch-year adoption of the three CEX goods

	Cellular phone service (1994)	Home internet service (1996)	Streaming video (2008)
Launch-year adoption rate (pp)	3.65	6.30	1.09
<i>Effect of ten years of age</i>			
pp, no controls	−0.40 (0.09)***	−0.79 (0.09)***	−0.23 (0.05)***
% of the adoption rate, no controls	−11.1%	−12.5%	−21.6%
pp, controls	−0.38 (0.12)**	−0.51 (0.12)***	−0.13 (0.05)*
% of the adoption rate, controls	−10.5%	−8.1%	−11.5%
R^2 , no controls	0.003	0.012	0.006
R^2 , controls	0.033	0.081	0.012
Consumer units	11,191	13,951	15,296
Adopters	423	936	168

Weighted linear probability models on the consumer units interviewed in each good's launch year; the outcome is positive spending on the good in any observed month of that year, in percentage points. Ten years of age in percentage points and as a share of the adoption rate, without and with controls (income quintile, family size, children, education, home ownership) and dummies for the months observed; HC1 standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.