Advertising Business Cycles

Rohan Shah[†]

Current draft: 10th January, 2026

Latest version available here

Abstract

Advertising has become more effective over time. How has this affected the long-run economy and business cycles? I develop a dynamic stochastic general equilibrium model in which heterogeneous multi-product firms invest in advertising and in Research and Development (R&D). Advertising increases demand for existing products while R&D creates new products. I find that historical improvements in advertising effectiveness have increased output and consumption by between 1.4% and 5.8% due to a reallocation of resources from low productivity to high productivity firms. I also find that increasing advertising effectiveness speeds up the economy's recovery from a recession and reduces the volatility of the business cycle. Moreover, advertising and R&D are substitutes or complements depending on the part of the business cycle in which the economy finds itself. These results highlight the importance of advertising on the aggregate economy.

Key words: R&D, Advertising, Economic Recovery, Product Innovation, Intangible capital

JEL Codes: E32, E22, M37, O30

[†] The University of Mississippi, rpshah1@olemiss.edu

1. Introduction

Advertising has become more effective over time: in particular, the advent of the internet has enabled firms to target their advertisements to different groups of consumers much more easily.[1] Total advertising spending in the US has been very pro-cyclical with a correlation of 0.83 with GDP.[2] Given this high pro-cyclicality, to what extent has the increase in advertising effectiveness affected the aggregate economy and business cycles in particular?

The increase in effectiveness could mean that if firms reduce their advertising spending in response to an aggregate decline in the economy the severity of the recession is exacerbated as the reduction in advertising leads to further reductions in demand and output. For example, the correlation between US output and advertising spending increased from 0.52 during 1960-1980 to 0.87 during 2000-2020, indicating that the effect of advertising on output might have increased over time. Alternatively, the increase in the effectiveness of advertising could mean that firms are better able to ameliorate the effects of any negative aggregate shock by increasing their advertising spending only slightly to get a greater effect on their demand, such that the recession is not as severe as it would have been otherwise. Similar opposing mechanisms could also prevail in the economy's long-run steady-state.[3]

I examine the impact of the increase in advertising's effectiveness on the economy's long-run steady-state and business cycles using a dynamic stochastic general equilibrium model of heterogeneous firms. In my model, multi-product intermediate goods firms can choose how much to spend on advertising to add to their existing stock of advertising in the next period. Increases in the firm's own stock of advertising increase the final goods producer's demand for that firm's products. I also allow firms to invest in R&D to increase their chances of developing a new product in the next period to add to their existing stock of products. This reflects the idea, espoused since Fogg-Meade (1901), that new products and advertising are closely associated and Baslandze et al. (2023)'s finding

[1] The ultimate measure of the effectiveness of advertising is how much extra revenue (and profit) is obtained from an extra unit spent on advertising (i.e. the elasticity of profit with respect to advertising), with more effective advertising resulting in a higher elasticity as the same amount of advertising spending results in higher profits. However, in practice, this is difficult to measure, so proxies such as the cost of advertising or the ability for advertising to be targeted are used instead (both of these proxies imply a higher impact of the same advertising spend on profits). For further discussion on advertising effectiveness and the implications of the rise of targeted internet advertising see, for example, Goldfarb (2014), Qi (2019), Baslandze et al. (2023), and Cavenaile et al. (2023).

[2] For more evidence on the pro-cyclical nature of advertising see Srinivasan, Lilien and Sridhar (2011); Hall (2012); and Molinari and Turino (2018).

[3] For example, Gijsenberg et al. (2009) find that advertising becomes less effective during recessions, whereas Tellis and Tellis (2009) find that the decline in aggregate advertising during a recession means that an individual firm's advertising might stand out more and therefore be more effective during a recession. Steenkamp and Fang (2011) also find that advertising is more effective during a recession.

that higher advertising effectiveness and the number of products is closely linked.[4] This approach also lets me examine the extent to which advertising and R&D spending are substitutes or complements: firms can choose to spend on either advertising or R&D and might not be able to afford spending on both (such that they are substitutes), but higher advertising can also increase the returns to R&D through increasing the demand for a new product and vice versa (such that they could be complements).[5]

I find that increases in the effectiveness of advertising have resulted in higher output and consumption in the long-run: output is between 1.4% and 5.8% higher and consumption is 3.1% and 5.6% higher under current levels of advertising effectiveness compared to historical levels. Higher advertising effectiveness means that less productive firms spend less on R&D and more on advertising, while more productive firms spend more on both R&D and advertising. Not only does this indicate that the substitutability / complementary nature of advertising and R&D can differ across firms, it means that the less productive firms increase their output through raising the demand for their pre-existing products while more productive firms increase their output by both increasing the demand for their products and developing more new products. Both of these factors serve to increase total output and consumption in the long run.

Over the business cycle the effect of increased advertising effectiveness depends on how that increased effectiveness is implemented. If it happens via a reduction in the cost of advertising, then a recession is much shorter as firms are better able to afford to use advertising as a means to ameliorate the impact of the downturn on demand for their products in the first few periods of the recession. When advertising is more effective firms switch from R&D to spending more on advertising (indicating that the two are substitutes) during the initial periods of the recovery. This dampens the negative impact of the recession. After a few periods the increase in the stock of advertising means that the returns to developing new products are higher such that firms spend more on both advertising and R&D and the recovery is rapid (i.e. advertising and R&D are now complements), with output actually overshooting the steady-state level after 8-10 years and then returning back down to the steady-state level over time. This rapid recovery is not possible when advertising costs more because firms cannot afford to increase their advertising spending when the shock hits, which means that the subsequent increase in R&D and stock of products also does not happen. Instead, when the cost of advertising is higher, the economy takes roughly 50 years before output returns up to its steady-state level.

[4] In particular, Fogg-Meade (1901) states that "when [a firm] wishes to introduce new goods, [the firm] must make them attractive. [The firm] must excite desire by appealing to imagination and emotion." and "[w]hen a new article is to be put on the market ... advertising expenses are always calculated among the initial costs."

[5] The importance of allowing for this heterogeneity given the data that some firms treat R&D and advertising to be complements and others treat them as substitutes is discussed in detail in Appendix A.

On the other hand, when increased advertising effectiveness happens via an increase in the impact of advertising on demand for a firm's product it serves to reduce the volatility of the economy. In particular, firms do not need to increase their advertising spending by as much in order to obtain the same amelioration effect and, as such, do not need to cut R&D by as much. Moreover, when advertising is more effective this smaller increase in advertising means that output does not overshoot its steady-state level as much (as there is a smaller increase in the stock of advertising, R&D spending, and the stock of products). Indeed, when advertising has less of an effect on a firm's demand, output overshoots the steady-state level by 0.25% compared to only 0.2% (relative to the steady-state) when advertising is more effective. As such, increases in advertising effectiveness helps to increase long-run output, speed up the economy's recovery from a recession, and dampen the volatility of the economy.

I contribute to three main strands of literature. First, I contribute to the literature concerning the effect of advertising on the aggregate economy. The closest work to mine is Molinari and Turino (2018), which uses a dynamic stochastic general equilibrium model with a representative firm that produces a single product and can advertise. They examine how advertising affects the business cycle and find that the incorporation of advertising leads to higher levels and volatility of consumption. In contrast, I model heterogeneous multi-product firms that can also invest in R&D to capture the interplay between a firm's R&D and advertising decisions and examine how changes in the effectiveness of advertising affect long-run aggregates and business cycles. This heterogeneity means that I can account for the empirical fact that some firms treat advertising and R&D as substitutes and others treat them as complements rather than modelling them just as complements as in Molinari and Turino (2018). As a result, I find a more muted impact of advertising effectiveness on long-run aggregates as well as substantial differences in how increases in advertising effectiveness affect business cycles (depending on how the change in the effectiveness of advertising is modelled).

The incorporation of R&D to examine its interplay with advertising is also examined by Cavenaile and Roldan-Blanco (2021) albeit using productivity-focused R&D rather than the product-focused R&D in my model. They analyse the interaction between this R&D and advertising to determine how it affects firm growth. Their finding that advertising and R&D are substitutes drives their result that reductions in advertising effectiveness results in higher long-run output as firms switch away from advertising to R&D. In contrast, and consistent with empirical evidence, I find that R&D and advertising can be complements as advertising increases the return to developing a new product (whereas advertising achieves the same effect as R&D in their model). Similarly, Cavenaile et al. (2023) focuses on household awareness resulting from advertising in a representative agent model and finds that information frictions result in welfare losses; my results

are consistent with these as I find that more effective advertising leads to increases in consumer welfare.

My results also reflect those of Dinlersoz and Yorukoglu (2012) which focuses on how firms grow over time when advertising provides informational value to consumers. In particular, my model includes the trade-off they identify between 1) a direct effect of higher advertising effectiveness leading to a firm's own advertising increasing demand for their product by a larger amount; and 2) an indirect effect of increased advertising effectiveness leading to lots more advertising in general such that an individual firm's advertising has to compete more with other firms' advertising. In addition, my model's findings match Afrouzi, Drenik and Kim (2023)'s empirically derived result that more effective advertising means more productive firms do more advertising and have higher output. However, in my model this effect arises due to these firms' higher advertising acting as a complement to their R&D such that they develop more new products, increasing output (rather than via an effect on mark-ups as in their model).

I also contribute to the literature examining the interplay between firm advertising and R&D. In contrast to most work in this area that focuses on R&D leading to productivity improvements, I focus on R&D that results in the introduction of new products. This reflects the importance of the interplay between new products and advertising highlighted by Argente et al. (2021) and Argente, Lee and Moreira (2024)'s results that young firms and firms with new products can grow more quickly by spending more on advertising.

My results complement those in Cavenaile and Roldan-Blanco (2021) and Cavenaile et al. (2024): each of those studies develops an endogenous growth model in which firms can invest in advertising to increase demand for their product and / or spend on advertising to increase their productivity. As advertising and R&D have the same effect (increasing demand for the firm's product) in their model, they find that R&D and advertising are substitutes. In contrast, I allow R&D and advertising to have different effects, with the result that whether or not R&D and advertising are substitutes depends on the firm and / or the part of the business cycle in which the economy finds itself. Nonetheless, similar to my results, they find that increases in the effectiveness of advertising leads to higher output albeit via a different mechanism.

Conversely, Ignaszak and Sedlacek (2023) find that advertising and R&D are complements when it comes to firm growth even when R&D improves a firm's productivity. They use a model in which firms must first acquire customers in order to sell their products. This approach differs from mine in that I allow firms to still make sales even if they have no advertising (reflecting the data that a large proportion of firms make sales despite not doing any advertising), and as mentioned above, I find that the substitutability / complementarity of advertising and R&D can differ across firms and

over time.[6] Indeed, this varying nature of the substitutability / complementarity of advertising and R&D is also found by Askenazy, Breda and Irac (2015). However, in their framework it is the competitive structure of the economy that determines whether or not advertising and R&D are substitutes or complements whereas in mine it depends on the firm and the part of the business cycle in which the economy is.

Finally, I also contribute to the literature examining the increasing importance of intangible capital in the modern economy. Haskel and Westlake (2018) provide an overview as to what distinguishes intangible capital from tangible capital as well as presenting evidence regarding firms' increased spending on intangibles such as R&D and advertising. Furthermore, Mcgrattan and Prescott (2014) show that adding intangible capital to a representative agent real business cycle can help to explain the increase in labour productivity despite a fall in output and hours worked in 2008 and 2009, while Crouzet and Eberly (2019) link the rise in the importance of intangible capital to changes in the level of concentration in the US economy. Ridder (2024) formally models the mechanism underlying how the rise of intangibles can help to explain the slowdown of productivity growth and the rise of market power. In contrast, I specifically focus on the impact of improvements to the effectiveness of advertising and how it interacts with R&D (instead of grouping all forms of intangible together), which allows me to examine how different types of intangible capital interact with each other both in the long-run and over the business cycle.

The rest of this paper proceeds as follows. Section 2 presents my model of multi-product firms and advertising, while Section 3 describes how I calibrate and solve the model non-linearly. Section 4 presents an analysis of the effectiveness of advertising on the economy's long-run steady-state and Section 5 shows how the change in advertising's effectiveness affects business cycle fluctuations and the recovery of the economy from a simulated recession. Section 6 concludes and briefly describes some potential avenues for future work.

2. Model

In this Section I set out a discrete-time dynamic stochastic general equilibrium model and characterise agents' optimisation problems, adapting Shah (2024) to incorporate advertising by intermediate goods firms. There are three categories of agents: a representative household that lives forever, a final goods producer that is perfectly competitive, and a mass of monopolistically-competitive heterogeneous multi-product intermediate goods producers. The representative household obtains utility from

[6] The fact that many firms do not do any advertising yet still make sales is also found by Kaiser (2005) for German firms that introduce new products.

consumption of the single consumption good C and its leisure time $1 - L$. The household owns the firms and receives all profits.[7] The household's problem is stated formally in Section 2.1.

The final goods producer takes the output of each product from the intermediate goods producers and aggregates them into the single consumption good using a CES aggregator with an elasticity of substitution θ between each good. The final goods producer's demand for each product from the intermediate goods producers depends on each firm's stock of advertising a relative to the total aggregate stock of advertising A with an effect that depends on an exponent α (a higher value of α means that advertising has a larger effect on the final goods producer's demand for a particular good, keeping each firm's stock of advertising and the aggregate stock of advertising constant). The price the final goods producer pays for each intermediate good is p. Note that a firm's stock of advertising applies to all of the products it owns, reflecting the concept of "umbrella advertising" (or brand advertising) as examined by Wernerfelt (1988). I present the problem of the final goods producer in Section 2.2.

Each multi-product intermediate goods producer is characterised by its endogenous measure of products N that only it can produce, its endogenous stock of advertising a, its persistent idiosyncratic exogenous research productivity ε_r, its persistent idiosyncratic exogenous advertising productivity ε_a, and its persistent idiosyncratic exogenous output productivity ε_o. Hence, there is a distribution μ of intermediate goods firms over these four characteristics.[8] The fact that firms have different productivities for each of output, research, and advertising in my model captures the fact that, for example, some firms are more effective at conducting research than are other firms, and some firms are more effective at conducting advertising campaigns than are other.[9]

A firm's stock of advertising depreciates by δ_a each period. Moreover, the intangibility of a firm's stock of advertising means that a firm cannot reduce its stock of advertising other than through letting it depreciate (i.e. a firm cannot sell off, or have negative investment in, its advertising stock).[10] This implies that an advertising investment made in one period automatically has an effect in future periods as well (i.e. that there is

[7] Households also have access to a full set of state-contingent claims, but as the household is representative in my model, these are in zero net supply in equilibrium. As such, I do not explicitly model these claims here for convenience.

[8] Each persistent idiosyncratic shock is an independent continuous process that follows an AR(1) process with a zero mean (in logs), persistence ρ_o and a disturbance standard error of σ_o for the output shock; persistence ρ_a and a disturbance standard error of σ_a for the advertising productivity shock; and persistence ρ_r and a disturbance standard error of σ_r for the research productivity shock.

[9] An example of the latter would be John Lewis: each new John Lewis advertising campaign, particularly their Christmas campaign, garners a lot of media attention in newspapers and television news reports above and beyond the advertising "spots" that are purchased for the campaign itself.

[10] The treatment of advertising as an intangible asset follows the likes of Haskel and Westlake (2018) and Crouzet and Eberly (2019).

a "stock of advertising" that depreciates) and is highlighted by Dehez and Jacquemin (1975)'s finding that "goodwill" arises from past advertising expenditures and affects affects current sales. It is also noted in Hall (2012)'s statement that "advertising is an investment that has a lasting benefit, extending beyond the period of the expenditure itself".[11]

The amount of each intermediate good that is produced depends on the stock of advertising and idiosyncratic output productivity of the intermediate goods firms that owns that product. As will be shown in Section 2.2, this means that the price p of a particular intermediate good is a function of an intermediate goods firm's advertising and output productivity. The intermediate goods firm's problem is presented in Section 2.3.

My method of modelling advertising allows three different approaches to investigating the impact of changes in the effectiveness of advertising. In one approach, I can simply alter the cost of investment in advertising I_a that is required to obtain an extra unit of advertising next period a'. Under this approach, higher advertising effectiveness is represented by a lower cost of investment required for the same amount of additional advertising next period. Alternatively, I can alter the exponent α on advertising in the CES aggregator with a decrease in the value of this exponent meaning that advertising has a lower impact on the demand for an intermediate goods firm's products, thereby representing a reduction in advertising's effectiveness. As such, I am able to examine three different versions of a change to the effectiveness of advertising: 1) changing just the cost of advertising; 2) changing just the exponent on advertising; and 3) changing both the cost and the exponent on advertising.

Incumbent intermediate goods firms face a probability π_d of exogenous exit each period. They produce output y of each product using labour ℓ. Firms that did not receive the exogenous exit shock choose their investment in advertising for the next period I_a and how much R&D investment R to make in order to gain the possibility of innovating a new product. This probability of innovation is a function $\Phi(R)$ (which is increasing in R and bounded between 0 and 1).[12] If an incumbent firm wishes to make an investment in R&D, it must pay a fixed cost ξ_r before doing so.[13]

[11] For further evidence regarding the treatment of advertising as a stock see, for example, Argente et al. (2021); Cavenaile et al. (2023); Afrouzi, Drenik and Kim (2023); and Steenkamp and Fang (2011).

[12] This endogenously varying probability of innovating a new product is what necessitates my use of a heterogeneous firm framework. Even absent any idiosyncratic output and research productivity across firms, each period some firms would be successful at innovating a new product and some firms would not be successful. This naturally creates a distribution of firms across the number of products they own, thereby creating firms that are heterogeneous.

[13] Hence, the R&D costs in this model are in terms of units of the consumption good. This follows the likes of Akcigit and Kerr (2018),Clementi and Palazzo (2016), Dekle et al. (2021), Hamano and Zanetti (2017), Klette and Kortum (2004), and Minniti and Turino (2013) in this area. An alternative formulation

My implementation of R&D potentially resulting in the introduction of a new product (rather than, for example, an improvement in productivity or quality) captures the association of new product introductions and advertising found by Argente et al. (2021) and Argente, Lee and Moreira (2024). Moreover, my distinct treatments of R&D and advertising (R&D as a period-by-period expenditure and advertising as a stock) follows the likes of Dinlersoz and Yorukoglu (2012), Cavenaile and Roldan-Blanco (2021), Ignaszak and Sedlacek (2023), Cavenaile et al. (2024) and others. A firm's R&D expenditure benefits from a subsidy τ_r from the government that is funded via a lump sum transfer T from households.[14] Note that a successful innovation only affects a firm's stock of products; it does not affect a firm's exogenous idiosyncratic idiosyncratic productivities of output, advertising.

An incumbent firm's stock of products depreciates exogenously at rate δ_N each period. This reflects the lifecycle of individual products, whereby a new product makes most of its revenues (and profits) for a firm in the first few years of its introduction.[15]

Each period there is a π_d mass of new entrants with one product and a stock of advertising ψA (i.e. ψ times the aggregate stock of advertising) that enters to replace the exogenously exiting firms. Note that the exogenous mass of entrants also pay the fixed cost of R&D before entering to reflect the fact that they would have had to conduct some R&D to develop a new product before entering. Finally, firms are subject to aggregate uncertainty in the form of an aggregate productivity shock z that varies over time (transition probabilities given by π_z) and affects the productivity of intermediate goods producers.

2.1. Households

An infinitely-lived representative household obtains utility from the single consumption good C and leisure $1 - L$ (where L is the household's supply of labour), and has a subjective rate of time preference β. For each unit of labour it supplies it earns wages $w(z, \mu)$. The household owns shares in the firms, as represented by the vector λ, and pays lump sum taxes T to the government.

Let the ex-dividend price of shares in a firm with products N, stock of advertising a, output productivity ε_o, research productivity ε_r, and advertising productivity ε_a be

would scale the R&D costs by the prevailing real wage (an approach taken by, for example, Aghion and Howitt (1992) andAyres and Raveendranathan (2023)). The impact of this modelling choice are outside the scope of this paper.

[14] The inclusion of an R&D subsidy captures the fact that R&D and advertising are subject to different fiscal incentives. In particular, firms do not receive any subsidy (or tax credit) for advertising but they do receive one for R&D, affecting the incentives firms face when deciding how much to spend on each (and potentially affecting the trade-off between the two if a firm considers them to be substitutes).

[15] See, for example, Argente et al. (2021), which shows that new products lose about 50% of their sales (on average) within two or three years of their introduction, with products only making about 30% of their initial sales as little as six years after their introduction.

$\rho_x(N', a', \varepsilon_o', \varepsilon_r', \varepsilon_a'; z', \mu)$; and the share price including dividend be $\rho_i(N, a, \varepsilon_o, \varepsilon_r, \varepsilon_a; z, \mu)$. Households maximise lifetime discounted utility via their choice of consumption, labour, and share-holdings. In other words, households maximise Equation (1) subject to the household budget constraint Equation (2)

$$V^h(\lambda; z, \mu) = \max_{C,L,\lambda'} U(C, 1-L) + \beta E_{\pi_z} V^h(\lambda'; z', \mu') \tag{1}$$

$$\begin{aligned} C + \int \rho_x(N', a', \varepsilon_o', \varepsilon_r', \varepsilon_a'; z', \mu)\lambda'(d[N' \times a' \times \varepsilon'_o \times \varepsilon'_r \times \varepsilon'_a)] \\ = w(z, \mu)L + T + \int \rho_i(N, a, \varepsilon_o, \varepsilon_r, \varepsilon_a; z, \mu)\lambda(d[N \times a \times \varepsilon_o \times \varepsilon_r \times \varepsilon_a) \end{aligned} \tag{2}$$

Following, Khan and Thomas (2013) households take the real wage and dividend-inclusive value of their current shares as given and choose their current consumption, hours worked and the number of new shares priced ex-dividend. The household's labour-leisure decision determines the equilibrium real wage w such that $w(z, \mu) = \frac{D_2 U(C,1-L)}{D_1 U(C,1-L)}$. In addition, the household optimal choice of shares establishes the firm's stochastic discount factor such that I can incorporate that stochastic factor in a more parsimonious way in the firm problem itself. In other words the firm problem is written with any current period profit scaled by $q(z, \mu) = D_1 U(C, 1-L)$.

I use Rogerson-Hansen preferences such that $U(C, 1-L) = \log C + B(1-L)$ where B is the household's disutility of working.

2.2. *Final goods producers*

The perfectly competitive final goods producer uses a CES aggregator to produce the single aggregate output good $Y(z, \mu)$ according to equation (3).

$$Y(z, \mu) = \left[\int^{\bar{N}} (1 + \frac{a}{A})^{\alpha} y(\varepsilon_o, z, \mu)^{\frac{\theta-1}{\theta}} dn \right]^{\frac{\theta}{\theta-1}} \tag{3}$$

It maximises profits by choosing how much output $y(\varepsilon_o, z, \mu)$ of each of the N intermediate products produced by an intermediate goods firm with advertising stock a and idiosyncratic output productivity ε_o to purchase at price $p(a, \varepsilon_o, z, \mu)$ in order to make aggregate output $Y(z, \mu)$, with $\bar{N}$ representing the total number of products in the economy and A representing the total stock of advertising. [16]

[16] N is the measure of products owned by an individual firm while $\bar{N}$ is the total across the entire economy. In other words, $\bar{N}$ is obtained by summing up N over all firms in the economy. Hence, summing over all products present in the economy is the equivalent to first summing the number of products each firm has and then adding up the number of firms weighted by their number of products.

The term $(1 + \frac{a}{A})^{\alpha}$ represents how an intermediate goods firm's advertising affects the demand for its product. In particular, holding the aggregate stock of advertising A constant, an individual firm with a higher stock of advertising gets a larger weight in the final goods producer's CES aggregator and therefore receives a higher demand for its products (with this higher demand also being reflected in a higher price as shown below). In this way, an intermediate goods firm's stock of advertising can be thought of as affecting consumer (or the final goods firm's) "perception" of the intermediate good, as per Cavenaile and Roldan-Blanco (2021) and Dixit and Norman (1978) - an increase in (relative) advertising $\frac{a}{A}$ means that there is an increase in demand from the final goods producer for the intermediate firms' product(s).[17] The exponent α determines how much a change in an intermediate firm's (relative) stock of advertising affects the final goods producer's demand for its products, which a higher value of the exponent resulting in advertising having a larger impact on demand (i.e. a higher exponent means that advertising is more effective).

My approach assumes that any advertising conducted by an intermediate goods firm affects demand at the "brand" (i.e. firm) level. This reflects Erdem and Sun (2002)'s finding that even product-specific advertising has effects at the brand level (i.e. that there are positive spillovers for an entire brand even if only one product is advertised) as well as Kaiser (2005)'s result that customers are often already aware of a firm's entire range of products through advertising on a subset of specific products. This also mirrors Cavenaile and Roldan-Blanco (2021)'s approach to including firm-level advertising, as well as Wernerfelt (1988)'s modelling of "umbrella advertising".

Moreover, this approach also captures both the direct and indirect effects of advertising highlighted by Dixit and Norman (1978) and Dinlersoz and Yorukoglu (2012). In particular, $(1 + \frac{a}{A})^{\alpha}$ is increasing in an individual firm's stock of advertising a such that there is a positive direct effect on the price and / or demand for an intermediate goods firm's products from an increase in advertising. However, this weight is decreasing in the aggregate stock of advertising A such that there is a negative indirect effect (as higher aggregate advertising reduces the impact any increase in an individual firm's advertising has) if all firms increase their advertising.[18]

My incorporation of firms having a baseline level of sales even if they do not have any stock of advertising (the unit value at the start of the weight that results from advertising), I am able to capture the large proportion of firms that do not make any advertising

[17] This set-up also captures the idea in Afrouzi, Drenik and Kim (2023) and Ignaszak and Sedlacek (2023) that firms spend on advertising to increase their customer base (although the exact mechanism through which advertising does this is not modelled here, potential mechanisms have been explored by Dinlersoz and Yorukoglu (2012) and Sedlacek and Sterk (2017)).

[18] This also captures Cavenaile et al. (2024)'s finding that firms can compete in advertising such that rival firms increasing their advertising reduces the impact that an individual firm's own advertising can have.

spending at all. Without the unit value, a firm would not make any sales unless it had some non-zero stock of advertising, and therefore all firms in the model would have to undertake at least some advertising.

Using this set-up, the final goods producer solves equation (4) subject to equation (3).

$$\max_{Y,y} P(z,\mu)Y(z,\mu) - \int^{\bar{N}} p(a,\varepsilon_o,z,\mu)y(\varepsilon_o,z,\mu) \tag{4}$$

Solving this maximisation problem and normalising the price $P(z,\mu)$ of the single aggregate output good to 1 means that the price of each intermediate good for a firm with stock of advertising a and idiosyncratic output productivity ε_o is given by equation (5).

$$p(a,\varepsilon_o,z,\mu) = (1+\frac{a}{A})^{\alpha}\left(\frac{Y(z,\mu)}{y(\varepsilon_o,z,\mu)}\right)^{\frac{1}{\theta}} \tag{5}$$

This price of an intermediate good depends on both the output of the intermediate good and the level of aggregate output (as well as the aggregate and firm's own stock of advertising). However, this price itself affects how much of each individual product an intermediate goods firms produces. As such, there is a feedback loop between aggregate output and the output of an intermediate product that means that firms need to forecast current aggregate output when deciding how much to produce within a particular period. This requires a forecasting rule for output within the current period, alongside the forecasting rule for current-period consumption / marginal utility when solving the model (this is discussed in more detail in Section 3.2).

The fact that the idiosyncratic output productivity ε_o and the stock of advertising a differ over intermediate goods firms means that the price of intermediate goods p will differ across firms, but the proportional effect of changes in aggregate output Y and z will predominantly be the same across all firms (such that changes in z will result in very similar deviations of the price per product from its steady-state value across all firms). In other words, different values of ε_o and a simply shift the level of the price per product and do not substantially affect its response to changes in z and Y, with firms that have more advertising or a lower exogenous idiosyncratic output productivity having a higher price for their product.

2.3. Intermediate goods producers

At the start of a period each intermediate goods producer finds out their persistent idiosyncratic research, advertising, and output productivities, and whether or not they have been hit with the exogenous death shock. Firms also learn the level of the aggregate

productivity shock at this time. High values of research and advertising productivity means that the same amount of investment in, respectively, R&D and advertising can be made more cheaply for those firms than it can for low levels of those productivities.

After observing these shocks, the intermediate goods firms hire labour and produce output of each of their own products using a Cobb-Douglas production function with weight γ on labour (where γ is between zero and one such that there are decreasing returns to the input). A firm's stock of advertising also affects its output through the effect on the price of the intermediate product as described above in Section 2.2. The profit for each individual intermediate good is given by $\pi(a, \varepsilon_o, z, \mu) = \max_{\ell,p} p(a, \varepsilon_o, z, \mu) y(\varepsilon_o, z, \mu) - w(z, \mu)\ell$ subject to $p(a, \varepsilon_o, z, \mu) = (1 + \frac{a}{A})^{\alpha} \left(\frac{Y(z,\mu)}{y(\varepsilon_o, z, \mu)}\right)^{\frac{1}{\theta}}$ where $y(\varepsilon_o, z, \mu) = z\varepsilon_o \ell^{\gamma}$ is the output of that individual product.[19] Note that each product owned by a firm is produced independently of any other product owned by a firm such that the production function of an individual product does not depend on the number of other products a firm owns.[20]

After production has occurred, those firms that did not receive the exogenous death shock choose how much to invest in advertising and in R&D. Firms (including those that received the exogenous death shock) then pay out any profits as dividends to the household. Each period there is also a π_d mass of potential entrants (with idiosyncratic output, advertising, and research productivities according to the ergodic distribution of each shock) that enter with advertising stock ϕA and pay ξ_r to do so. Therefore, the timing within a period is as follows:

1. Incumbents find out their idiosyncratic research, advertising, and output productivities for the current period, and whether or not they are subject to the exogenous death shock.
2. All incumbents hire labour and produce output
3. Incumbents continuing onto the next period choose how much to invest in advertising and in the possibility of innovating a new product.
   (a) Firms that are subject to the exogenous exit shock simply discard all of their stock of advertising and their stock of products (they are unable to sell them off).
   (b) All incumbent firms then pay any dividends to the household.
   (c) The π_d mass of entrants each pay the fixed cost of R&D and pay for the ψA stock of advertising with which they will start the next period.

[19] This means that a firms revenue (in terms of its inputs) is $(1 + \frac{a}{A})^{\alpha} Y(z, \mu)^{\frac{1}{\theta}} (z\varepsilon_o)^{\frac{\theta-1}{\theta}} l^{\gamma \frac{\theta-1}{\theta}}$ such that a firm's returns to scale with respect to inputs (labour and stock of advertising) is roughly given by $\alpha + \gamma \frac{\theta-1}{\theta}$.

[20] Moreover, since a firm's research productivity only affects the effectiveness of its R&D in developing new products, a firm's research productivity has no effect on the output of products that already exist.

More formally, let W represent the value of an incumbent firm at the start of the period and W^1 be the value of a firm that does not receive the exogenous death shock. The incumbent firm's optimisation problem can be defined recursively as

$$\begin{aligned} W(N, a, \varepsilon_o, \varepsilon_a, \varepsilon_r; z, \mu) = \pi_d \left[\max_{\ell, p} q(z, \mu)[N\pi(a, \varepsilon_o, z, \mu)] \right] \\ + (1 - \pi_d) W^1(N, a, \varepsilon_o, \varepsilon_a, \varepsilon_r; z, \mu) \end{aligned} \tag{6}$$

Equation (6) states that the firm takes the probability of an exogenous death into account. If the firm receives the exogenous death shock it chooses its labour to maximise profits across all of its products and pays out dividends to the household.[21] If the firm does not receive the exogenous death shock it hires labour to maximise output, then chooses its R&D and advertising investments to maximise W^1, which is given below:

$$\begin{aligned} W^1(N, a, \varepsilon_o, \varepsilon_a, \varepsilon_r; z, \mu) = \max_{\ell, p, R, k'} q(z, \mu)[N\pi(a, \varepsilon_o, z, \mu) - \frac{(1 - \tau_r)R}{\varepsilon_r} - \mathbb{I}_r \xi_r - \frac{I_a}{\varepsilon_a}] \\ + \beta \Phi(R) E_{\pi_z} E_{\pi_{\varepsilon_r}} E_{\pi_{\varepsilon_a}} E_{\pi_{\varepsilon_o}} W((1 - \delta_N)N + 1, a', \varepsilon_o', \varepsilon_a', \varepsilon_r'; z', \mu') \\ + \beta (1 - \Phi(R)) E_{\pi_z} E_{\pi_{\varepsilon_r}} E_{\pi_{\varepsilon_a}} E_{\pi_{\varepsilon_o}} W((1 - \delta_N)N, a', \varepsilon_o', \varepsilon_a', \varepsilon_r'; z', \mu') \end{aligned} \tag{7}$$

subject to

$$\begin{aligned} R &\geq 0 \\ I_a &\geq 0 \\ 0 &\leq N\pi(a, \varepsilon_o, z, \mu)_i - \frac{(1 - \tau_r)R}{\varepsilon_r} - \mathbb{I}_r \xi_r - \frac{I_a}{\varepsilon_a} \\ a' &= (1 - \delta_a)a + I_a \\ p(a, \varepsilon_o, z, \mu) &= (1 + \frac{a}{A})^{\alpha} \left(\frac{Y(z, \mu)}{y(\varepsilon_o, z, \mu)} \right)^{\frac{1}{\theta}} \\ \mu' &= \Gamma(\mu) \end{aligned} \tag{8}$$

where $\Phi(R)$ represents a firm's probability of innovating a new product given their investment in R&D (R) and τ_r is the R&D subsidy.

The firm must choose its investment in advertising I_a today, without knowing whether it will be successful in innovating a new product. As such, there some interplay between a firm's decision to invest in R&D and advertising. The constraints mean a firm cannot

[21] Note that a firm makes a profit of $\pi(a, \varepsilon_o, z, \mu)$ per product and has N products, such that its total operating profits from all of its products is simply $N\pi(a, \varepsilon_o, z, \mu)$. This means that a firm is indifferent between making a profit on one of its products compared to another one of its products. In other words, there is no cannibalisation of sales (or operating profits) within a firm's own portfolio of goods beyond that which occurs via the aggregate output CES function.

make negative investments in R&D and it cannot equity finance any R&D or advertising investments (i.e. it must always pay non-negative dividends).

This formulation allows me to capture the fact that advertising can either: 1) be a complement for R&D by increasing the returns to innovating as a firm will obtain more profit for each new product it innovates via advertising increasing the price of the product; or 2) be a substitute for R&D as firms cannot pay negative dividends (via the constraint $0 \leq N\pi(a, \varepsilon_o, z, \mu)_i - \frac{(1-\tau_r)R}{\varepsilon_r} - \mathbb{I}_r\xi_r - \frac{I_a}{\varepsilon_a}$) and therefore have to choose between spending more on R&D or more on advertising if they are unable to afford both without violating this constraint. The importance of allowing for this heterogeneity given the data that some firms treat R&D and advertising to be complements and others treat them as substitutes is discussed in detail in Appendix A. Moreover, the absence of any term $(1-\delta_a)a$ in the firm's dividends and the presence of the constraint that $I_a \geq 0$, reflect the idea that a firm's stock of advertising is intangible and therefore cannot be sold off.

Note that a successful innovation (with probability $\Phi(R)$) results in a firm having $(1-\delta_N)N+1$ products next period.[22] However, with probability $1-\Phi(R)$ a firm does not make a successful innovation and has only $(1-\delta_N)N$ products next period. This means that a firm's choice of R&D investment R affects the number of products that it will have next period, such that this number of products is an endogenous individual state variable. I use a logit function for this probability given an investment in R&D. In other words, $\Phi(R) = \frac{e^{\chi R}}{1+e^{\chi R}}$ where the parameter χ will be calibrated to match target moments. $\mathbb{I}_r$ is an indicator function taking the value zero when a firm makes zero investment in R&D and a value of one when it makes a positive R&D investment; hence $\mathbb{I}_r\xi_r$ represents the fixed cost of being able to carry out some non-zero amount of R&D investment. Finally, $\mu' = \Gamma(\mu)$ represents the law of motion of the distribution of firms.

2.4. Competitive equilibrium

Let $\mathbb{I}^O$ represent those firms in the distribution that are continuing incumbent firms and $\mathbb{I}^{\mathbb{N}}$ denote the π_d mass of entrants, a Recursive Competitive Equilibrium is a set of prices p, w, q, ρ_i, ρ_x, and functions for quantities ℓ, R, I_a, A, Y, C, L, λ', T and values V^h, W, W^1, that solve the firm and household problems and clear markets as described below:

1. W and W^1 solve Equations (6) and (7), and ℓ, R, and I_a are the associated policy functions for incumbent firms

[22] Assuming that a firm can only add one product (at most) per year allows me to capture Luttmer (2011)'s result that small firms grow much quicker relative to large firms: in my model a one-product firm adding one product doubles their stock of product (and, hence, increases their sales and employment by a lot) whereas a, say, five-product firm adding one product increases their number of products by 20% and so increases their sales and employment by much less relatively than does the smaller firm.

2. V^h solves Equation (1) with C, L, λ being the household's associated policy functions
3. The distribution of firms over $N \times a \times \varepsilon_o \times \varepsilon_a \times \varepsilon_r$ is given by μ and evolves according to $\mu' = \Gamma(\mu)$
4. Aggregate output is given by $Y = \left[\int^{\bar{N}} (1 + \frac{a}{A})^{\alpha} y^{\frac{\theta-1}{\theta}} dn\right]^{\frac{\theta}{\theta-1}}$
5. The labour market clears $L = \int N\ell(N, a, \varepsilon_o, \varepsilon_a, \varepsilon_r; z, \mu)\mu(d[N \times a \times \varepsilon_o \times \varepsilon_a \times \varepsilon_r)$
6. The goods market clears $C = Y - \int \left(\mathbb{I}^O[\frac{R}{\varepsilon_r} + \mathbb{I}_r \xi_r + \frac{I_a}{\varepsilon_a}] + \mathbb{I}^{\mathbb{N}}[\mathbb{I}_r \xi_r + \frac{\psi A}{\varepsilon_a}] \right) \mu(d[N \times a \times \varepsilon_o \times \varepsilon_a \times \varepsilon_r])$
7. The aggregate stock of advertising is given by $A = \int a\mu(d[N \times a \times \varepsilon_o \times \varepsilon_a \times \varepsilon_r)$
8. Lump sum taxes equal the total amount of R&D subsidies $T = \int \mathbb{I}^O \tau_r R\mu(d[N \times a \times \varepsilon_o \times \varepsilon_a \times \varepsilon_r])$
9. Prices ρ_i and ρ_x clear the market for shares

where N is the measure of products owned by an individual firm while $\bar{N}$ is the total across the entire economy.[23]

3. Quantitative Strategy

In this Section, I set out how I calibrate the values of parameters in my model to match target moments. I first describe the calibration of the model's non-stochastic steady-state before setting out how the parameters governing the exogenous aggregate shocks are determined.[24]

3.1. Steady-state Calibration

Solving the (non-stochastic) steady-state of the model requires iterating on guesses of aggregate consumption (which affects the wage), output (which affects the price intermediate goods firms charge for their output), and the aggregate stock of advertising (which affects the returns to an individual firm's stock of advertising) until all three converge. Conditional on these guesses of consumption, output, and aggregate advertising, the first step involves obtaining the optimum (non-negative) choice of a firm's stock of advertising given a choice of R&D investment assuming that a firm can afford this optimum level. I then find the optimum choice of R&D assuming that a firm is able to afford both the optimum choice of R&D and the implied optimum choice of advertising investment.

[23] $\bar{N}$ is obtained by summing up N over all firms in the economy. Summing over all products present in the economy is the equivalent to first summing the number of products each firm has and then adding up the number of firms weighted by their number of products.

[24] The specific approach and algorithms I use to solve my model are set out in Appendix B.

For the next step, I check if the firm can actually afford these unconstrained optimum choices without violating the condition that they must pay non-negative dividends. If a firm can afford these unconstrained optimums then that is what they choose. If choosing the unconstrained optimums would imply the firm paying negative dividends (i.e. having to equity finance) then I re-solve the firm's optimum choice of R&D and advertising. Specifically, I find a constrained firm's optimum choice of R&D ensuring that such a firm's choice of advertising investment does not result in a violation of the non-negative dividend constraint.

Using these optimum choices of R&D and advertising investment, I iterate on the distribution of firms until it converges and then find the resulting new values of consumption, output, and aggregate stock of advertising. These new values are used to update the original guesses of consumption and output using Broyden's algorithm until convergence. More details regarding this solution algorithm are presented in Appendix B.1.

I use this solution method to calibrate the steady-state of my model against an overidentified set of target moments. One period in my model is equivalent to one year; this is to allow for consistency between the moments related to the introduction of new products and advertising, which are only available at the annual level. Table 1 shows the model and targeted steady-state moments that relate to aspects of the model other than advertising and the distribution of firms over products.[25] Although all parameters are determined jointly the rows of the table show which parameters are most-closely associated with each target. The labour share is determined by γ the exponent on labour in the intermediate goods firms' Cobb-Douglas production function, while the household disutility of working B targets the proportion of time the household spends working. The elasticity of substitution θ in the final goods producers' CES aggregator function is set to target the midpoint between the 1976 and 2006 mark-ups found by Loecker, Eeckhout and Unger (2020).

The proportion of output spent on R&D investment, and the proportion of firms that make zero R&D investment each year, are most closely affected by the persistence ρ_r and volatility σ_r that govern the idiosyncratic research productivity shock ε_r and the fixed cost of making a positive investment in R&D ξ_r. Finally, the entry rate is determined by the exogenous chance of death π_d while the relative average size of entrants compared to incumbents and the proportion of total employment by entrants are determined by the proportion ψ of the aggregate stock of advertising with which entrants enter. My

[25] Compustat data are used for two targets and although those data only include public firms (a small subset of the total population of firms in the US) those firms account for a large proportion of private R&D spending (roughly 90%) and advertising spending (roughly 60%) as per Cavenaile and Roldan-Blanco (2021).

model performs very well in matching each of these aggregate targets. For example, it almost exactly matches standard aggregate moments such as the labour share and hours worked, as well as the proportion of output that is spent on R&D. It also matches the relative importance of new entrants compared to incumbent firms

Table 1: Model's steady-state moments for standard aggregate measures

Moment	Main parameter	Source	Model	Data
Labour Share	Cobb-Douglas Exponent γ	Cooley & Prescott (1995)	62%	60%
% of time worked	Household disutility of working B	Literature	32%	33%
Aggregate mark-up	Elasticity of substitution θ	De Loecker, Eeckhout & Unger (2020)	36.25%	36.25%
R&D investment as % of output	Idiosyncratic shock terms for ε_r, fixed cost of positive R&D investment ξ_r	NSF	1.8%	1.8%
Proportion of firms with 0 R&D investment	Idiosyncratic shock terms for ε_r, fixed cost of positive R&D investment ξ_r	Compustat 1980-2020	51.5%	53.1%
Entry rate	Exogenous death probability π_d	BDS	10.6%	10.6%
Relative average size of entrants vs incumbents	Entrants' starting proportion of aggregate advertising stock ψ	BDS	29.0%	25.6%
Proportion of total employment by entrants	Entrants' starting proportion of aggregate advertising stock ψ	BDS	3.3%	3.0%

Note: The proportion of firms with 0 R&D investment is obtained from Compustat annual data over the period 1980-2020 having dropped all financial and regulated firms from the sample. Outlier observations in which a firms sales, assets, employees, "property, plant, and equipment", depreciation, accumulated depreciation, or capital expenditures are recorded as being negative are also dropped from the sample.

Turning to the target moments that relate to the distribution of firms over number of products and introduction of new products), these are shown in Table 2 and are governed predominantly by the logit probability parameter χ, the product depreciation rate δ_N, and the fixed cost of making a positive investment in R&D ξ_r, with the targets for these moments taken from Bernard, Redding and Schott (2010). My model captures the distribution of firms over R&D expenditure and the introduction of new products, with the only major discrepancy in that area being a too high proportion of large firms that add a new product.[26] Despite this latter aspect, my model gets very close to matching

[26] Despite not being explicitly targeted, my model gets close to Broda and Weinstein (2010)'s finding that roughly 25% of products in any one year are "new": in my model, this number can be obtained by multiplying the steady-state number of continuing firms (0.894) by the average probability of an incumbent or entrant firm developing a new product (51%) and adding that to the replacement 0.106 entrants that start producing to get 0.56 new products introduced each year. This is then divided by the total number of products in the economy (this is just the average number of products per firm (1.85) as there is a unit mass of firms at all times). As such, the proportion of products that are new in any one year is (0.894*0.51+0.106)/(1.85) = 30.3%.

the proportion of output by large firms that is accounted for by those large firms that add a new product and therefore still captures the relative importance of large firms adding a new product in the aggregate.

Table 2: Model's steady-state moments for the distribution of firms over number of products

Moment	Model	Data
Mean number of products per firm	1.85	1.975
% of firms that are single-product	59.5%	61%
Proportion of firms that add a new product	51%	39%
Proportion of output accounted for by firms adding a product	88%	78%
Proportion of output accounted for by multi-product firms	92%	87%
Proportion of multi-product firms adding a product	89%	68%
Proportion of multi-product firm output accounted for by multi-product firms adding a product	89%	87%
Proportion of large firms adding a product	89%	44%
Proportion of large firm output accounted for by large firms adding a product	89%	80%

Note: Although all moments in my model are determined jointly by all parameters, the specific parameters that predominantly determine the moments shown in this table are the logit probability parameter χ, the product depreciation rate δ_N, and the fixed cost of making a positive investment in R&D ξ_r. The targets for all of these moments are obtained from Bernard, Redding and Schott (2010)

Finally, the target moments related to the importance of advertising in the economy are shown in Table 3. In particular, each of advertising investment as a proportion of output, advertising spending as a proportion of firm sales, and the proportion of firms with zero advertising investment are predominantly determined by the exponent α on advertising in the CES aggregator, the depreciation δ_a of the advertising stock each period, and the persistence ρ_a and volatility of disturbances σ_a of the idiosyncratic advertising productivity shock. Crucially, for the purposes of assessing the importance of changes to the effectiveness of advertising, my model does relatively well in matching the targeted advertising moments. In particular, it has slightly too much advertising relative to output, but slightly too little advertising relative to firm's sales, but still gets relatively close to matching both targets and therefore captures the relative importance of advertising in the economy. Moreover, by almost exactly matching the proportion of firms that have zero advertising expenditure, my model is able to capture this aspect of the distribution of advertising investment across firms.

Overall, therefore, my model performs well in matching this over-identified set of targets. Table 4 shows the parameter values chosen for my model (all expressed in levels). As the period of my model is one year, the household rate of time preference β is set to 0.96 to obtain a risk-free real interest rate of 4%, while the R&D subsidy τ_r is set at 5% as per the

Table 3: Model's steady-state moments for advertising targets

Moment	Source	Model	Data
Advertising spending as % of output	Baslandze et al (2023)	2.8%	2.2%
Advertising spending as % of firm sales	Cavenaile & Blanco (2021)	2.6%	4%
Proportion of firms with 0 advertising expenditure	Compustat 1980-2020	64.5%	64.1%

Note: Although all moments in my model are determined jointly by all parameters, the specific parameters that predominantly determine the moments shown in this table are the exponent on advertising in the CES aggregator α, the depreciation of the advertising stock δ_a, and the persistence ρ_a and volatility of disturbances ε_a to the idiosyncratic advertising productivity shock. The proportion of firms with 0 advertising expenditure is obtained from Compustat annual data over the period 1980-2020 having dropped all financial and regulated firms from the sample. Outlier observations in which a firms sales, assets, employees, "property, plant, and equipment", depreciation, accumulated depreciation, or capital expenditures are recorded as being negative are also dropped from the sample.

level of the subsidy in the USA reported by OECD (2022); these are the only parameters that are set externally. Note that the depreciation rate of advertising stock at roughly 12% per year is of the same order of magnitude as Steenkamp and Fang (2011)'s finding that about 85% of the effect of a firm's advertising on its profit carries over from the previous year (i.e. the depreciation in that effect is roughly 15%).

To validate my model, I examine its performance against a number of untargeted moments. Specifically, I assess the model's performance in capturing the firm-size distribution; the growth paths of new firms' R&D and advertising expenditures; and the relationship between a firm's advertising expenditures and its sales. My model performs well at capturing the observed patterns in the data even in these untargeted moments. The results of this validation exercise are discussed in more detail in Appendix C

3.2. Calibration of aggregate shocks

To incorporate aggregate uncertainty into my model I assume that there are shocks to aggregate productivity z. These shocks are modelled as an AR(1) continuous process with a (log) zero mean, persistence ρ_z and a disturbance standard error of σ_z. As such, aggregate productivity is given by $\log z' = \rho_z \log z + \varepsilon_z$ where $\varepsilon_z \sim N(0, \sigma_z^2)$. I estimate the persistence ρ_z and standard deviation of the shock σ_z parameters via the Solow Residual obtained from Hodrick-Prescott filtered data on real annual US GDP and private capital as well as total civilian employment hours created by Cociuba, Prescott and Ueberfeldt (2018) over the period 1954-2007. This results in an estimated persistence $\rho_z = 0.744$ and standard deviation $\sigma_z = 0.0128$. I discretise the process across three gridpoints using the Rouwenhorst algorithm to obtain its values and transition probabilities.

Table 4: Parameter values

Parameter	Description	Value	Parameter	Description	Value
β	Household rate of time preference	0.96	ρ_r	Persistence of idiosyncratic research productivity	0.87
τ_r	R&D subsidy	0.05	v_r	Standard deviation of idiosyncratic research productivity shock	0.0085
γ	Cobb-Douglas Exponent on labour	0.71	ρ_o	Persistence of idiosyncratic output productivity	0.99
α	Exponent on advertising in CES aggregator	0.18	v_o	Standard deviation of idiosyncratic output productivity shock	0.021
B	Household disutility of working	2.25	ρ_a	Persistence of idiosyncratic output productivity	0.75
θ	Elasticity of substitution between products	3.759	v_a	Standard deviation of idiosyncratic output productivity shock	0.05
δ_N	Depreciation rate for intermediate products	0.189	δ_a	Depreciation rate for advertising	0.12
χ	Logit probability parameter	167.7	ψ	Entrants' starting proportion of aggregate stock of advertising	0.099
ξ_r	Fixed cost of making a positive R&D investment	0.13	π_d	Exogenous probability of death	0.106

I then solve my model using the (modified) Krusell-Smith algorithm as per Khan and Thomas (2013) and approximate the distribution of firms using the aggregate number of goods $\bar{N}$, the aggregate stock of advertising A, the lag of the aggregate number of goods $\bar{N}_{t-1}$, the lag of the aggregate productivity shock z_{t-1}, and the lag of the number of goods held by single-good firms.[27] More details regarding the solution algorithm are presented in Appendix B.2.[28]

To obtain the business cycle statistics presented in Table 5 I simulate my model for 1,500 periods and then discard the first 500 periods of the simulation. The first row of the table shows the mean of each series, while the second shows their coefficient of variation. The third row shows the volatility of each series relative to that of output. As one would expect, consumption is less volatile than output due to households' desire to smooth consumption.

[27] This means that I need four forecast rules: one each for the aggregate number of goods next period, the aggregate stock of advertising next period, and output and marginal utility in the current period. I do not need forecasting rules for the lagged variables as the value they will have in period $t+1$ are already known at time t.

[28] I repeat this procedure for each separate specification of the model with different advertising effectiveness for which I present results in Section 5.2.1.

Similarly, the number of products $\bar{N}$ is more volatile than output even though R&D spending substantially less volatile than output when looking at relative standard deviations. Note, however, that the coefficient of variation (i.e. once the means of the respective series are taken into account) are very similar between output and the aggregate number of products. Also using the coefficient of variation, R&D and advertising expenditure are more volatile than the number of products and output, in line with what we would expect and consistent with Molinari and Turino (2018). Moreover, the aggregate stock of advertising A is less volatile than output due to the fact that firms cannot sell off their stock of advertising in the face of adverse shocks but can only let it decrease via depreciation.

Furthermore, as mentioned in Section 1, both advertising and R&D expenditures are pro-cyclical: R&D spending has a correlation of 0.94 with output in the model compared to 0.38 in the data, while advertising spending has a correlation with output of 0.55 in the model compared to 0.83 in the data. These untargeted moments are also consistent with the pro-cyclicality of advertising as found by Molinari and Turino (2018). More details regarding my model's ability to capture the path of the US economy over the business cycle (such as almost exactly matching the ratio of standard deviations of each of advertising and R&D spending with output) despite not explicitly targeting any such moments are presented in Appendix D.

Table 5: Simulated business cycle moments for baseline model

	Y	C	$\bar{N}$	Total R&D Spending	A	Total Advertising Spending
Mean	0.79	0.68	1.89	0.014	0.12	0.024
Coefficient of Variation	2.7%	2.7%	2.8%	4.0%	1.5%	3.3%
$\frac{\sigma_X}{\sigma_Y}$	(0.021)	0.87	2.52	0.03	0.08	0.04
Correlation with GDP	1.00	1.00	0.76	0.94	0.77	0.55
Auto-correlation	0.88	0.88	0.98	0.96	0.89	0.09

Note: These moments are obtained from a simulation of 1,500 periods of the full model, dropping the first 500 periods of that simulation. These moments have not been filtered or smoothed. The first two rows show, respectively, the mean and coefficient of variation of each series. The third row shows the standard deviation of each series relative to that of output, except for the first column, which shows the actual standard deviation of output. The fourth row shows each series' contemporaneous correlation with output, while the last row contains each series 1-period auto-correlation. All numbers are expressed in levels.

4. Impact of the effectiveness of advertising on the long-run steady-state

In this Section, I present an analysis of how advertising effectiveness affects the economy's non-stochastic long-run steady-state. As mentioned above, I model a change in the effectiveness of advertising in three ways: 1) just increasing the cost of advertising; 1) just reducing the exponent on advertising in the final goods producer CES aggregator;

and 3) increasing the cost and reducing the exponent at the same time. In all cases, the baseline version of the model is taken to represent the case when advertising has its current, modern, level of effectiveness and the alternative scenarios reflect the past (in which advertising was less effective). Hence, changing advertising's effectiveness means reducing its effectiveness either by increasing its cost or by reducing the exponent on advertising in the CES aggregator.

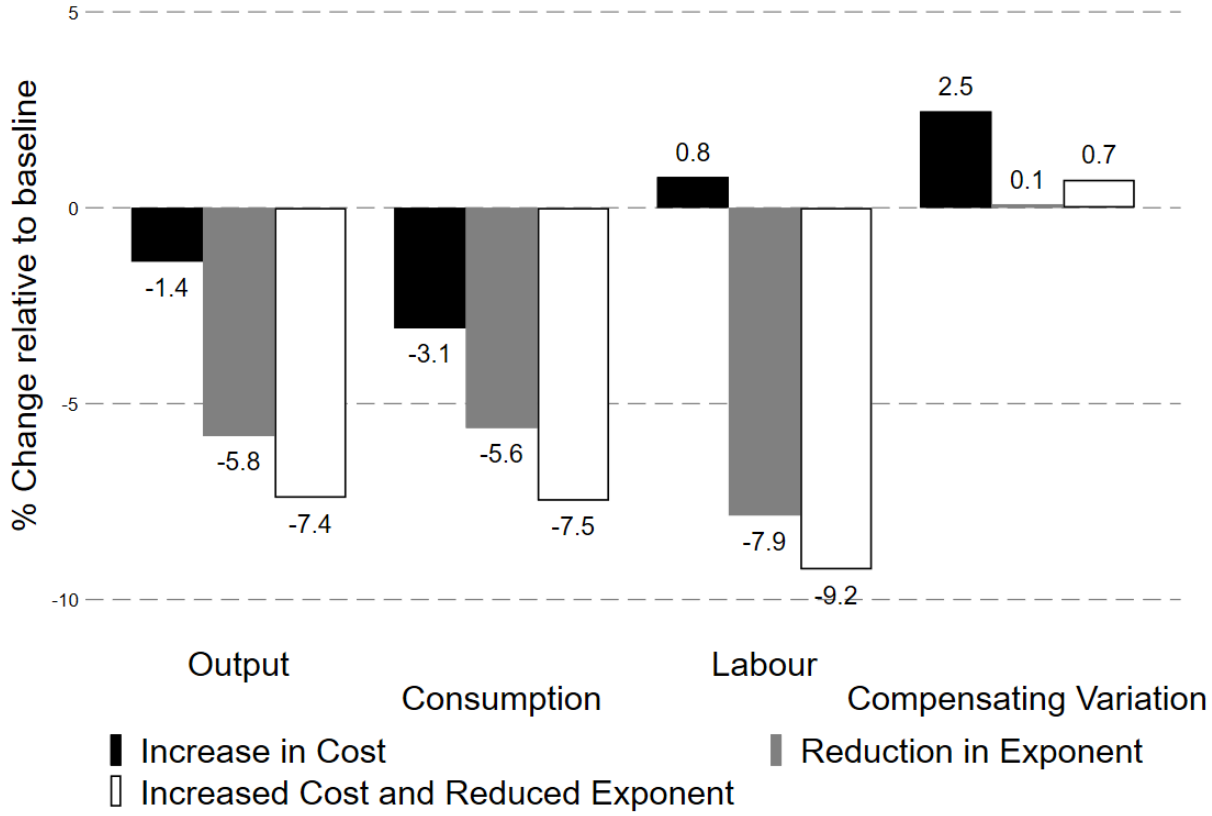
In the case of increasing the cost of advertising, that cost increased five times to reflect the fact noted in Cavenaile et al. (2023) that the implied cost per unit of (targeted) advertising has fallen by 80%; in other words, an additional unit of advertising stock next period now costs $5\frac{I_a}{\varepsilon_a}$ rather than just $\frac{I_a}{\varepsilon_a}$.[29] In the case of reducing the exponent α on advertising in the CES aggregator, I reduce that exponent on advertising to target a reduction in total advertising spending: advertising spending in 1980 was only 43% of what it was in 2020, so I reduce the value of α so that the steady-state advertising spending is only 43% of what it is in the baseline case (this requires changing the value of α to be 0.091, with all other parameter values unchanged). For the scenario in which both costs are increased (by the same amount as in the first scenario) and the exponent on advertising is reduced requires a further reduction in this exponent to match the same target such that α is 0.078 in that scenario.

Figure 1 shows the effect of reducing the effectiveness of advertising on the non-stochastic steady-state for output, consumption, labour, and a measure of welfare (the compensating variation). Each variable is expressed in terms of their change relative to the baseline specification of advertising effectiveness, with the black bars representing the change when costs are increased and the grey bars the change when the exponent on advertising in the CES aggregator is reduced. The white bars with a black border show the results when both the cost of advertising is increased and the exponent on advertising is reduced at the same time.

All methods of lowering the effectiveness of advertising reduce output and consumption. In the steady-state, a reduction in the effectiveness of advertising serves to reduce output, consumption, and household welfare. This is despite lower advertising effectiveness increasing aggregate R&D spending and the number of products. Specifically, the increase in the number of products is almost entirely from lower output productivity firms. In contrast, high output productivity firms actually reduce both their R&D and advertising spending such that their output is much lower and act as complements. As such, lower productivity firms treat R&D and advertising as substitutes in that they substitute away from advertising towards R&D when advertising is less effective. Hence, there is a difference in the interaction between R&D and advertising across firms.

[29] Cavenaile et al. (2023) also note that "the return to targeting goes from 0.048 in 2005 to 0.213 in 2014, nearly a five-fold increase": in other words, a nearly 80% reduction in the cost of advertising.

Figure 1: Comparison of Aggregate output, consumption, and labour in Steady-State by effectiveness of advertising



Note: The graph shows the impact of changing the effectiveness of advertising (adjusting by increasing the cost of advertising expenditure) for different aggregate variables, all expressed relative to their values in the baseline scenario. The black bars show the impact of reducing the effectiveness of advertising by increasing its cost while the grey bars show the impact of reducing advertising's effectiveness by decreasing the exponent on advertising. The white bars with a black border show the results when both the cost of advertising is increased and the exponent on advertising is reduced

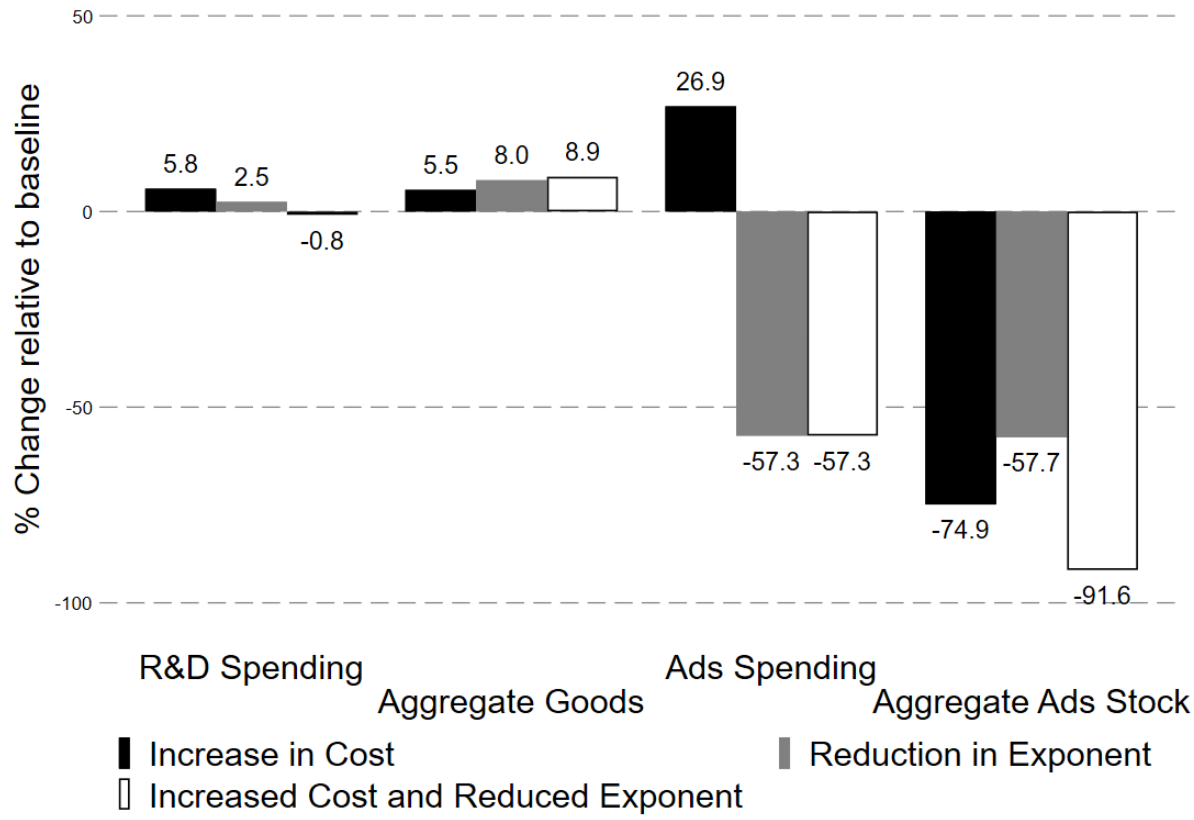
However, reducing the exponent on advertising reduces both output and consumption by much more than does increasing the cost of advertising. In addition, the combination of changing both at the same time generally increases the size of these effects. Note that the reduction in output when both costs are increased and the exponent is reduced is (slightly) larger than when each of those changes is made separately, although the opposite is true for consumption.

Moreover, although hours worked is reduced when the exponent on advertising is decreased, hours worked actually increases when the cost of advertising increases: this combination of changes to consumption and hours worked means that people are much worse off when advertising costs more, as shown by the bars for Compensating Variation (in which a positive number represents consumers being worse off): consumers would require 2.5% more income to compensate them for the reduction in consumption and increase in hours worked. However, consumers are only slightly worse off if advertising is less effective due to its impact on the CES aggregator being reduced (they require 0.1%

more income in this case). Combining both the increase in costs and the reduction in the exponent on advertising results in consumers being somewhere in between the two changes in effectiveness by themselves.

The reasons for these different effects of different methods of implementing changes in the effectiveness of advertising is shown by the changes' effects on aggregate R&D spending, number of goods, and advertising, each of which are shown in Figure 2. When examining aggregate effects the changes in these variables suggest that R&D and advertising are substitutes in the long-run: the aggregate stock of advertising is reduced and the aggregate number of goods are increased as a result of firms switching away from spending on advertising to spending on R&D instead. Note that even though advertising spending increases by nearly 27% (relative to the baseline) when the cost of advertising is reduced, the amount of "effective" advertising spending is much lower as total advertising spending would have had to increase by five times (i.e. by 400%) to completely offset the five times increase in cost of advertising.

Figure 2: Comparison of aggregate R&D, Number of goods, and Advertising in Steady-State by effectiveness of advertising



Note: The graph shows the impact of changing the effectiveness of advertising (adjusting by increasing the cost of advertising expenditure) for different aggregate variables, all expressed relative to their values in the baseline scenario. The black bars show the impact of reducing the effectiveness of advertising by increasing its cost while the grey bars show the impact of reducing advertising's effectiveness by decreasing the exponent on advertising. The white bars with a black border show the results when both the cost of advertising is increased and the exponent on advertising is reduced

Even though the increase in the number of goods is larger when the exponent on advertising is reduced compared to when the cost of advertising is increased, output falls by more in the former (as shown above). This is because reducing the effectiveness of advertising leads to firms with lower output productivity substantially increasing their R&D spending while high output productivity firms actually spend less on R&D. This means that the number of products and output of low productivity firms increase but the reduction in products for high productivity firms means that the decrease in output from these firms more than offsets the increased output of low productivity firms.

In other words, R&D and advertising tend to be substitutes for lower output productivity firms (as these two variables move in opposite directions) and complements for high output productivity firms (as advertising and R&D spending both fall when the effectiveness of advertising is reduced). This is due to the fact that lower output productivity firms are more likely to be bound by the non-negative dividend constraint (and thus must choose to spend on either R&D or advertising, with an increase in one necessitating a decrease in the other) whereas high output productivity firms are less likely to be constrained by having to pay non-negative dividends and, therefore, can increase (or decrease) both R&D and advertising spending at the same time. Hence, even though, in the aggregate, R&D and advertising might appear to be substitutes, for some firms they are complements in the steady-state. This effect is compounded when both advertising costs are increased and the exponent on advertising is reduced at the same time.

This means that the increase in the number of products is almost entirely from firms with low output productivity. Coupled with the large reduction in advertising stock across all firms, this means that total output (via the CES aggregator) is lower even though there are more products simply because it is the lower output productivity firms that have the largest increase in number of products. Moreover, this effect is stronger when the exponent on advertising is reduced compared to the costs of advertising being increased. This arises because the increase in costs and resulting reduction in aggregate stock of advertising actually encourages some smaller firms to increase their own stock of advertising (taking advantage in the reduced impact of the negative indirect effect) and substitute away from R&D such that the increase in the number of products is smaller. Conversely, all firms decrease their advertising when the exponent on advertising is reduced.

This can be seen when looking at the joint distribution of firms over their number of products and stock of advertising, as shown in Figure 3. The top panel shows this joint distribution for the baseline case of "modern-day" advertising effectiveness; the middle panel shows it for a the case in which advertising effectiveness is reduced through an increase in the cost of advertising; the third panel shows the joint distribution when

advertising effectiveness is decreased by reducing the exponent on advertising in the CES aggregator; and the bottom panel shows this for the combination of increasing costs and reducing the exponent on advertising. In each panel, the x axis is the firm's stock of advertising, the y axis is the firms stock of products, and the z axis is the proportion of firms that have that combination of advertising and products.

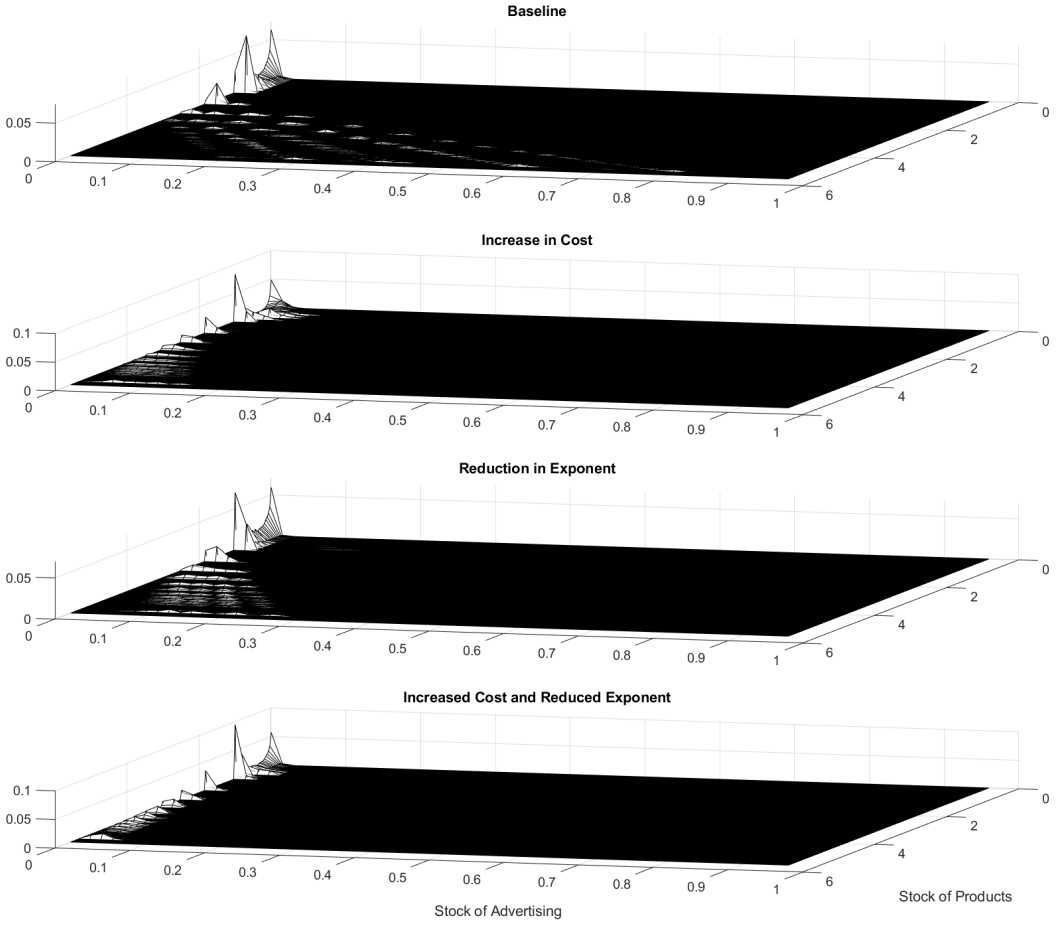
Focusing on the baseline case (the top panel) first, firms with a high number of products (roughly more than 2 or 3) have the most stock of advertising, with their stock of advertising tending to increase as their number of products also increases (with the variation in the peaks in this area of the graph reflecting the effect of different past and current combinations of the idiosyncratic productivity shocks). Firms with fewer than 2 or 3 products have very little advertising stock, other than a very small proportion of firms that have around 1 product and an advertising stock of roughly 0.15 (reflecting the exogenous new entrants).

As one would expect, reducing the effectiveness of advertising reduces the optimum stock of advertising for firms that already have a high stock of advertising: the amount of advertising for firms with more than 2 or 3 products is greatly reduced, albeit with that reduction more pronounced when the cost of advertising is increased (the second panel) compared to when the exponent on advertising is reduced (the third panel). This helps to explain why the aggregate advertising stock is much lower when the cost of advertising is increased: it stems from the large reduction in advertising stock held by these firms with a large number of products.

When costs are increased, there is a "new" mass of firms with roughly 0.075 stock of advertising and between 1 and 2 products, where there had been zero mass of firms in the baseline case. These firms are the ones for whom the large reduction in the stock of advertising (and, hence, the large decrease in the negative indirect effect of advertising) means it is now beneficial for them to have some stock of advertising themselves as it is more effective at boosting their revenues (as the positive direct effect of their own advertising remains unchanged).

However, comparing the baseline case to the case where advertising effectiveness is reduced by decreasing the exponent on the CES aggregator shows that there is little change in the stock of advertising held by these firms with between 1 and 2 products: in this case, although the negative indirect effect is reduced (as the aggregate advertising stock is lower), the positive direct effect of a firm's own stock of advertising is also reduced through the lower value of the exponent on advertising. Hence, when the exponent is reduced, these smaller firms have no incentive to increase their advertising, unlike when the cost of advertising is increased. This means that these smaller firms actually have a higher output when advertising costs are higher than they do when the exponent on advertising is reduced, such that aggregate output (and consumption)

Figure 3: Comparison of Distribution of firms over Products and Advertising by effectiveness of advertising



Note: The graph shows the impact of changing the effectiveness of advertising (adjusting by increasing the cost of advertising expenditure) on the joint distribution of firms over their stock of products and their stock of advertising. The top panel shows the baseline scenario, the second panel shows the effect of reducing the effectiveness of advertising by increasing its cost, the third panel shows the effect of decreasing the effectiveness of advertising by reducing the exponent on advertising; and the bottom panel shows the effect by increasing costs and decreasing the exponent at the same time. In each panel, the x axis is the firm's stock of advertising, the y axis is the firm's stock of products, and the z axis is the proportion of firms with that combination of products and advertising.

does not fall by as much when costs are increased compared to when the exponent on advertising is reduced. Moreover, the combined effect of increasing the cost of advertising and reducing its exponent almost entirely eliminates advertising by large and small firms alike (as seen in the bottom panel).

5. Impact of the effectiveness of advertising on the economy's recovery from a recession

In this Section I present results examining how changes in advertising effectiveness affect the aggregate economy's recovery from a recession. I first show the response of the

baseline model (i.e. current levels of advertising effectiveness) to a 1% drop in aggregate TFP that gradually returns to its steady-state level for the full model: for this analysis the economy is set to run for 150 periods at the median level of TFP before TFP drops by 1% in the 151st period and then gradually returns to its median value as governed by the persistence parameter ρ_z.

In the baseline case, I find that advertising provides firms with a means to ameliorate the drop in aggregate TFP such that output quickly returns to (and exceeds) its steady state-level within ten years of the initial drop in TFP. At the start of the recession firms switch away from R&D to advertising so as to benefit from this amelioration effect; this suggests that advertising and R&D are substitutes in the first few periods of a recession.

However, after a few periods, firms' stocks of advertising are sufficiently high that now there is a much larger benefit from having more products. Coupled with firms increasing their R&D investments, this drives the rapid recovery of output. During this time, R&D and advertising act as complements to each other. Subsequently, R&D and advertising remain above the steady-state level of a long time precisely because they complement each other and can only decrease through depreciation, with the result that the economy returns back down to its steady-state level after roughly 60 years (having surpassed the steady-state level roughly 10 years after the initial shock).

Next, I examine how reducing the effectiveness of advertising affects the economy's business cycle moments and its recovery from the same TFP shock as in the baseline case. In other words, I compare the "current" (baseline) case against the historical effectiveness of advertising. I do this for each of just increasing the cost of advertising; just reducing the exponent on advertising in the final goods producer CES aggregator; and increasing the cost and reducing the exponent at the same time. Over the business cycle, an increase in the effectiveness of advertising reduces volatility: output, consumption, the number of products, and other aggregate variables are all less volatile when advertising is more effective.

When it comes to the economy's recovery from a recession I find that lower advertising effectiveness increases the volatility of the recovery. In particular, when advertising is less affordable, firms are no longer able to ameliorate the impact of the initial drop in aggregate TFP by increasing their advertising spending. Hence, both R&D and advertising fall when advertising costs more, such that there is no uptick in R&D and the number of products a few periods into the recovery. This results in output taking a long time to start to increase after its trough, which the result that output only gets back up to the steady-state level after roughly 50 years.

In contrast, when advertising has less of an effect on the CES aggregator, firms over-compensate by increasing their advertising much more when the shock hits: they do this because they need even more advertising to ameliorate the initial drop in aggregate

TFP. However, this leads to such a large increase in the stock of advertising and number of products that output recovers even more quickly and ends up even higher above the steady-state level than it does in the baseline case. As such, a reduced impact of advertising on a firm's demand increases the volatility of output (and other aggregate variables).

These results are each discussed in more detail below.

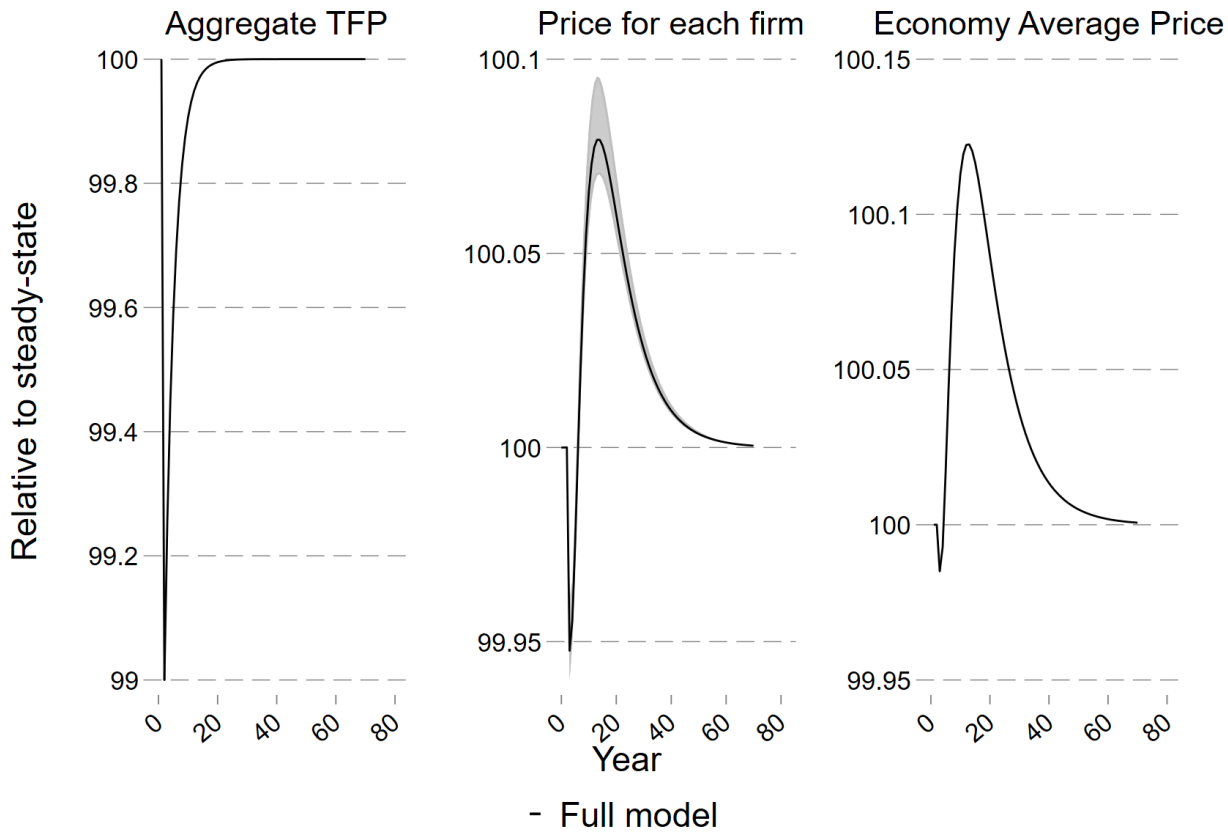
5.1. Aggregate Response in the baseline "present-day" scenario

At the outset it is useful to examine the path of prices for intermediate products in response to the shock to TFP, which is shown in Figure 4. The left panel of this figure shows the path of aggregate TFP over time. The middle panel shows (the distribution of) price paths for an intermediate product produced by a firm with advertising stock a and idiosyncratic productivity ε_o relative to the steady-state price for that particular type of firm. The grey shaded area in this panel shows the 10th and 90th percentiles of the distribution over prices for each firm, with a firm's price varying depending on their stock of advertising and output productivity according to $p(a, \varepsilon_o, z, \mu) = (1 + \frac{a}{A})^{\alpha} \left(\frac{Y(z,\mu)}{y(\varepsilon_o,z,\mu)}\right)^{\frac{1}{\theta}}$.

It is immediately clear that prices for products produced by firms with different combinations of advertising and output productivity follow a very similar path relative to their steady-state level. This path is one in which there is a slight delay in the price per product dropping after the TFP shock (due to aggregate output initially falling such that the ratio of aggregate output to firm output is level). This reduction in price is slightly more than 0.05%. Almost immediately after that, however, the price for each firm increases above the steady-state level: as discussed in more detail below, this is because firms invest heavily in advertising during those periods, such that this increase in advertising more than offsets the drop in aggregate productivity. This path of individual firm prices is also reflected in the average price of intermediate products for the economy as a whole, as shown in the right-hand panel.

The response of aggregate quantities to a 1% shock to TFP with a gradual return to its median level is shown in Figure 5, with each panel showing a different aggregate variable. The initial drop in aggregate TFP causes an immediate decrease in output and consumption, with consumption falling by more than output. This latter aspect is because firms increase their advertising spending substantially (by more than 1.5% compared to the steady-state) upon the shock hitting the economy. Despite the fact that they decrease their R&D spending slightly at the same time, this reduction is not as large as the increase in advertising spending, leading to consumption falling by more than output when the shock hits. Firms increase their advertising expenditure by such a large

Figure 4: Simulated path of price of intermediate products after a TFP recession in the baseline model



Note: The graphs show the simulated time path of the price of intermediate products in the full model economy. The economy is started at the steady-state distribution and runs for 150 periods with exogenous aggregate productivity at the median value of 1. In the 151st period aggregate TFP is subjected to a 1% reduction (i.e. is set to 0.99) before gradually recovering back to its median value of 1 at a rate determined by ρ_z. All series are expressed in terms of their percentage deviation from steady-state.

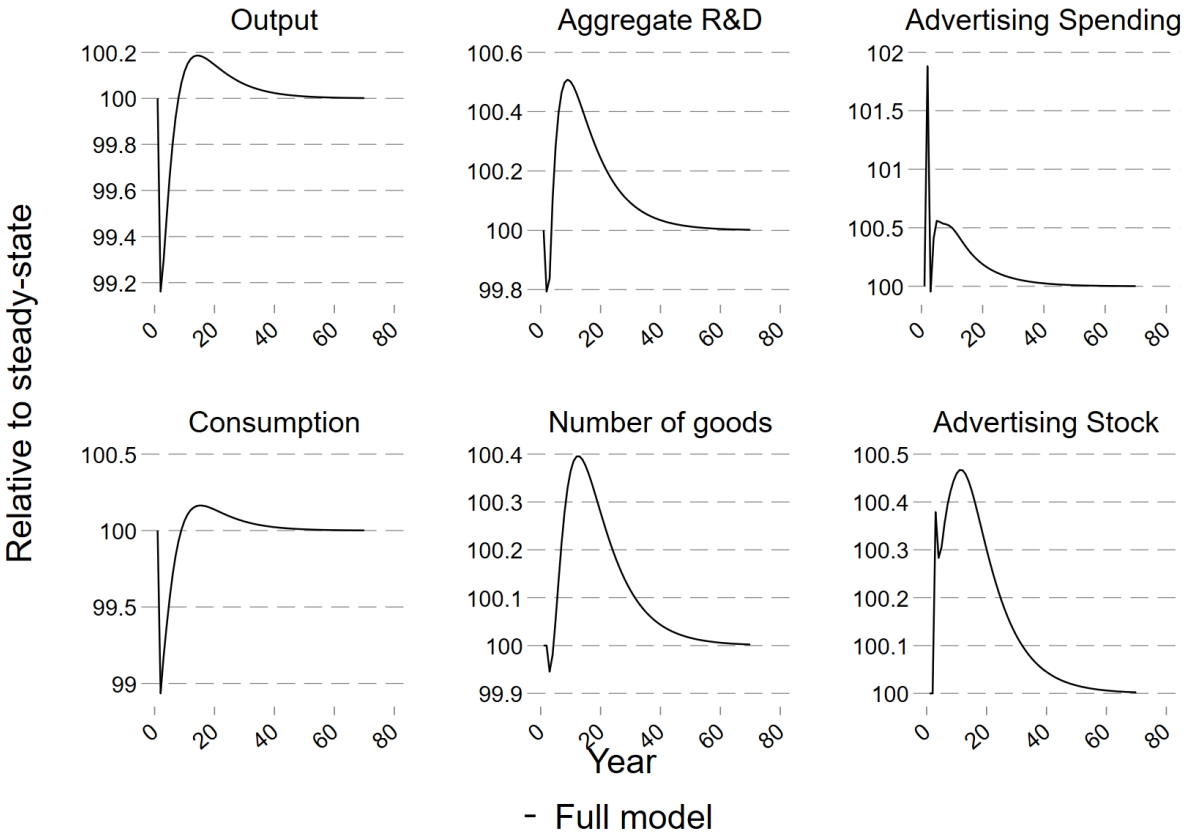
amount in order to try to ameliorate the effects of the reduction in aggregate TFP on their revenues.[30]

The fall in R&D is due to firms not being able to afford their previous levels of R&D at the same time as boosting their advertising spending and the returns to having new products are lower due to the drop in aggregate TFP. The fact that R&D and advertising spending move in opposite directions indicate that they are substitutes in the aggregate at this point in the economy's recovery. This is because the non-negative dividend constraint binds for more firms than prior to the shock (due to a reduction in revenues as a result of the shock to TFP) such that more firms now have to choose between either advertising or R&D, so firms that increase their advertising spending tend to have to reduce their R&D

[30] This increase in the aggregate stock of advertising in the face of a negative TFP shock is very different from the path of of (tangible) capital in a model that does not have advertising. For example, in both Khan and Thomas (2013) and Shah (2024), which incorporate physical capital, the stock of physical capital is always below its steady-state level (until the economy has fully recovered) when the economy has been hit by a negative TFP shock - i.e. at no point does the aggregate stock of physical capital exceed its steady-state level during the economy's recovery, in contrast to the path of the stock of (intangible) advertising in the model presented in this paper.

spending as a result. Firms switching from spending on R&D to spending on advertising at the start of a recession in my model is consistent with the patterns seen in the aggregate data, discussed in more detail in Appendix E (note, however, that even though R&D and advertising move in different directions at the start of the recession, they are positive correlated with each other across the entire recovery, also matching what we see in the data).

Figure 5: Simulated recovery of aggregate variables after a TFP recession in the full model



Note: The graphs show the simulated time path of aggregate variables in the full model economy. The economy is started at the steady-state distribution and runs for 150 periods with exogenous aggregate productivity at the median value of 1. In the 151st period aggregate TFP is subjected to a 1% reduction (i.e. is set to 0.99) before gradually recovering back to its median value of 1 at a rate determined by ρ_z. All series are expressed in terms of their percentage deviation from steady-state.

In the period just after the shock hits output and consumption already start to recover although the initial reduction in R&D spending has resulted in a slight fall in the number of products. More noticeably, the stock of advertising has increased substantially, resulting in firms over-correcting their advertising spending: the fact that aggregate TFP is still quite low relative to its median steady-state level and the stock of advertising has increased means the negative indirect effect of advertising via the presence of a higher denominator of $\frac{a}{A}$ results in advertising not being as attractive to firms. As such, firms reduce their advertising spending (so that it is roughly at the steady-state level during this period) and increase their R&D spending slightly instead. In other words, in the

first few periods after the shock hits, advertising spending is quite volatile but remains above the steady-state level. This is consistent with the path of US advertising spending in the years after the start of an NBER recession. This is discussed in more detail in Appendix F.

From the third period after the shock, the economy is less volatile. Now R&D and advertising spending both increase to roughly 0.5% above their steady-state levels. From this point, R&D and advertising serve as complements to each other. The higher R&D and advertising spending soon increase the aggregate number of products and stock of advertising, such that the returns to R&D and advertising positively reinforce each other. In other words, R&D and advertising are now complements at the aggregate level. This is due to the fact that TFP is now higher than it was at the start of the shock and the increase in advertising means that demand for firms' products is higher: both of these mean that fewer firms are now bound by the non-negative dividend constraint, such that firms can increase both R&D and advertising spending at the same time rather than having to choose between one or the other (as they did at the onset of the shock).

This positive reinforcement leads to rapid increases in output and consumption. In fact, both output and consumption surpass their steady-state levels roughly 8-10 years after the initial shock. Moreover, the positive reinforcement continues such that the number of products and stock of advertising peak roughly 15 years after the initial drop in TFP. This results in that output and consumption reaching their peak about 18 years after the initial shock.

By the time of this peak in output and consumption, aggregate TFP has returned very close to its steady-state level (and continues to get even closer to it), firms do not need to compensate for lower aggregate TFP by having higher advertising. In fact, the aggregate stock of advertising is now so high that the individual returns to having more advertising are lower via the indirect effect of advertising. This means that firms gradually reduce their advertising and R&D spending until they return to their steady-state levels after roughly 60 years. The fact this takes such a long time is due to: 1) both the number of goods and the stock of advertising only being able to decrease through depreciation (i.e. firms cannot sell off "unwanted" goods or advertising); and 2) advertising and R&D being complements. In other words, the slow reversions of the stock of advertising and the number of products to their steady-state level mean that the returns for a firm investing in R&D and / or advertising remain above their steady-state levels for a long time, such that firms continue investing in both of them above their steady-state levels until those returns dissipate as the stocks of goods and advertising get closer to their steady-state levels. This is also why output and consumption take nearly 60 years to return to their steady-state levels.

This recovery path makes clear that firms use advertising to try to ameliorate the effect

of the reduction in aggregate TFP: their initial attempts to increase their revenues through doing more advertising increase their (and the aggregate) stock of advertising (substituting away from R&D spending) such that their revenues do not fall as much as they would have absent this increase in advertising. After a few periods firm's stocks of advertising are such that there is now more benefit from having a higher number of goods, so firms increase their R&D spending back and advertising and R&D become complements. This then drives the rapid return (and eventual surpassing) of output and consumption to their steady-state levels.

5.2. Impact of advertising effectiveness on business cycles and economic recoveries

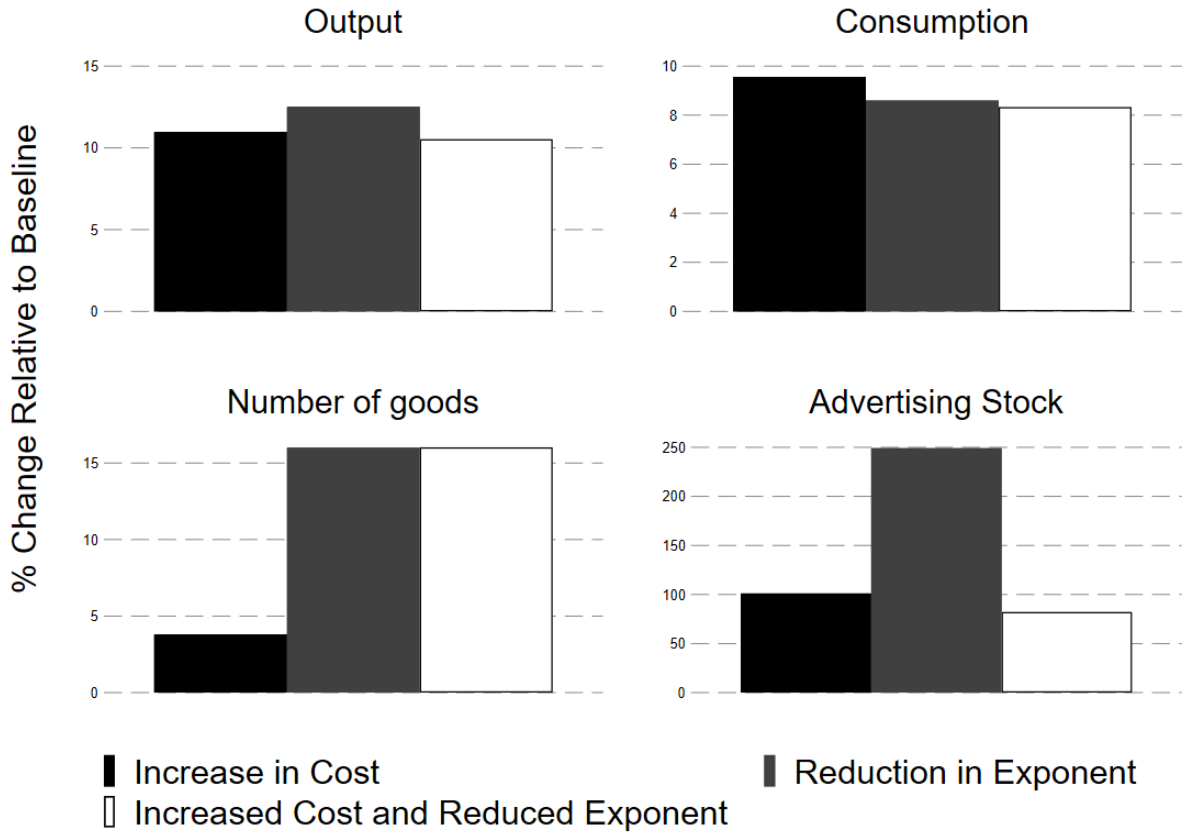
I analyse the impact of advertising effectiveness on the economy's business cycle fluctuations. First I show the the changes in volatility and correlation with output of various aggregate variables over a 1,500 period simulation of the economy (as was presented for the baseline case in Section 3.2) that result from changes in advertising effectiveness. This I examine how different levels of advertising effectiveness affect the economy's recovery from the same TFP-induced recession as was shown for the baseline case above.

5.2.1. Impact on business cycle moments

Figure 6 shows how the volatilities (as measured by the coefficient of variation) of aggregate series change when advertising is less effective. Each variable is expressed in terms of percentage change relative to the coefficient of variation in the baseline specification of advertising effectiveness. The black bars represent the change when costs are increased, the grey bars the change when the exponent on advertising in the CES aggregator is reduced, and the white bars with a black border the effect of both the cost of advertising increasing and the exponent on advertising decreasing. The implemented changes to parameters in each of these three scenarios are the same as described in Section 4.

It is immediately clear that reducing the effectiveness of advertising increases the volatility of the aggregate economy. For example, lower advertising effectiveness increases the volatility of output by roughly 10% and that of consumption by at least 8%. Moreover, the volatility of the stock of advertising is increased by at least 75% in each case, with a reduction in advertising's exponent in the CES aggregator leading to a 250% increase in the volatility of the advertising stock. In each case, the increased volatility is because firms cannot use advertising as much to ameliorate the effect of changes in aggregate TFP on the demand for (and price of) their products, thereby making their cashflow and R&D investments more volatile. This then feeds through into

Figure 6: Impact of the effectiveness of advertising on volatility of aggregate variables over the business cycle

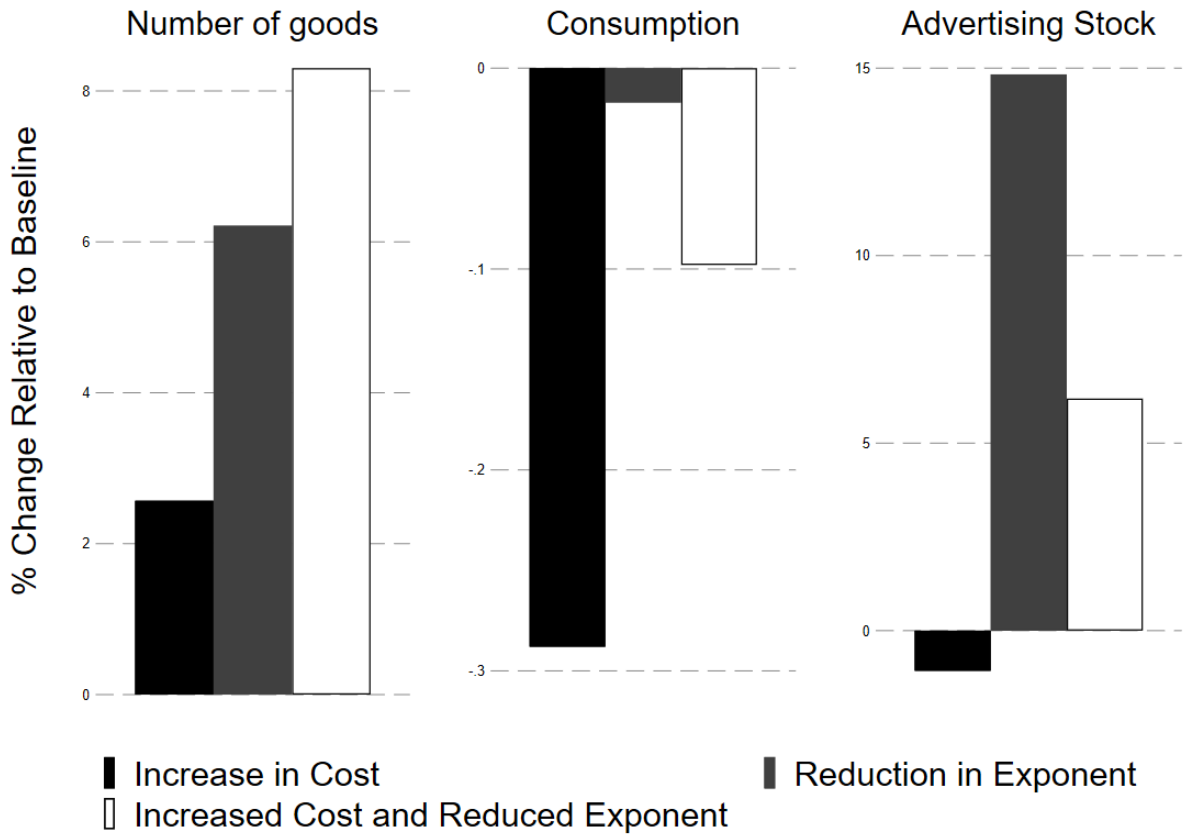


Note: The graph shows the impact of changing the effectiveness of advertising (adjusting by increasing the cost of advertising expenditure) for different aggregate variables, all expressed relative to their values in the baseline scenario. The black bars show the impact of reducing the effectiveness of advertising by increasing its cost while the grey bars show the impact of reducing advertising's effectiveness by decreasing the exponent on advertising. The white bars with a black border show the results when both the cost of advertising is increased and the exponent on advertising is reduced. The moments are obtained from a simulation of 1,500 periods of each full specification, dropping the first 500 periods of that simulation. These moments have not been filtered or smoothed.

increased volatility of R&D investments and the number of products, such that output and consumption are also more volatile.

Reducing the effectiveness of advertising also affects series' correlation with output, albeit to a much lesser extent than for their volatility, as shown in Figure 7. For example, reduced advertising effectiveness increases the correlation of the number of goods with output: this is because firms cannot reduce the impact of fluctuations in TFP by using advertising as effectively, such that the funds they have available to invest in R&D (and therefore develop more products) now depend much more on the level of output (which is greatly affected by aggregate TFP). This is also why the stock of advertising also tends to be more correlated with output. However, the effectiveness of advertising has little impact on the correlation between consumption and output: less effective advertising reduces the correlation between consumption and output by less than 0.3%.

Figure 7: Impact of the effectiveness of advertising on the correlation of aggregate variables with output over the business cycle

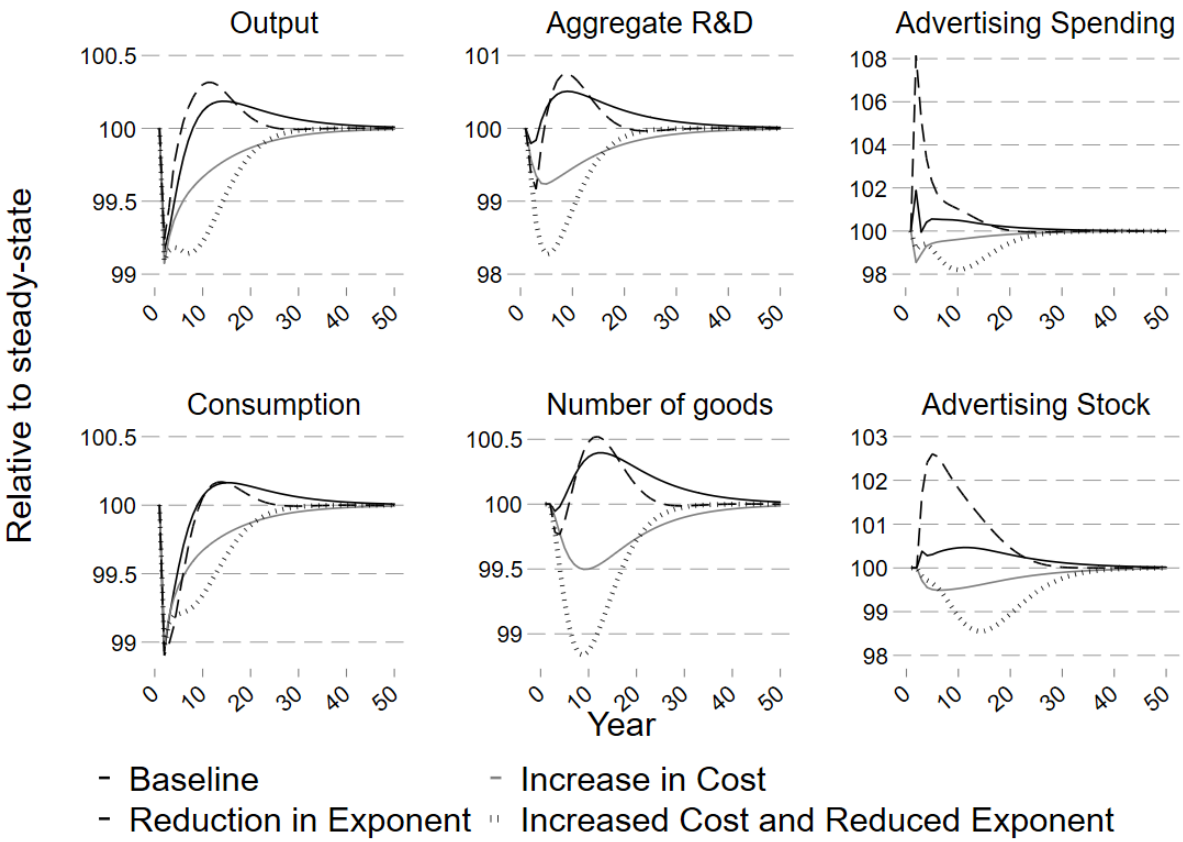


Note: The graph shows the impact of changing the effectiveness of advertising (adjusting by increasing the cost of advertising expenditure) for different aggregate variables, all expressed relative to their values in the baseline scenario. The black bars show the impact of reducing the effectiveness of advertising by increasing its cost while the grey bars show the impact of reducing advertising's effectiveness by decreasing the exponent on advertising. The white bars with a black border show the results when both the cost of advertising is increased and the exponent on advertising is reduced. The moments are obtained from a simulation of 1,500 periods of each full specification, dropping the first 500 periods of that simulation. These moments have not been filtered or smoothed.

5.2.2. Impact on the economy's recovery from a recession

Changes in the effectiveness of advertising also affect an economy's recovery from a recession. The values for the increase in costs and / or reduction in the exponent on advertising in the CES aggregator are the same for each specification as discussed in Section 4. Figure 8 shows the economy's recovery from a 1% drop in, and gradual return to steady-state of, aggregate TFP as was conducted for the baseline model in Section 5.1; the results from this baseline model are reproduced by the solid black line in each panel. The solid grey line shows the economy's recovery when advertising effectiveness is reduced via an increase in the cost of advertising, while the dashed black line shows the recovery when the effectiveness of advertising is reduced via a lower exponent on advertising in the CES aggregator. The dotted line shows the recovery when both the cost of advertising is increased and its CES exponent is reduced at the same time. Each panel shows the relevant variable relative to its steady-state in the particular model.

Figure 8: Simulated recovery of aggregate variables after a TFP recession when advertising effectiveness is changed



Note: The graphs show the simulated time path of aggregate variables in the full model economy. The economy is started at the steady-state distribution and runs for 150 periods with exogenous aggregate productivity at the median value of 1. In the 151st period aggregate TFP is subjected to a 1% reduction (i.e. is set to 0.99) before gradually recovering back to its median value of 1 at a rate determined by ρ_z. All series are expressed in terms of their percentage deviation from steady-state. The solid black line shows the response of the economy in the baseline model; the solid grey line shows the response for the economy in which the effectiveness of advertising has been reduced by increasing its cost; and the dashed black line shows the response for the economy in which advertising's effectiveness has been decreased by reducing the exponent on advertising. The dotted line shows the recovery in an economy in which both the cost of advertising has been increased and the exponent on advertising in the CES aggregator has been reduced at the same time.

First consider when the cost of advertising has increased (the solid grey line). In contrast to the baseline case in which firms increased their advertising spending upon the initial reduction in TFP to try to ameliorate the effects of the recession, a higher cost of advertising means that when the shock hits firms cannot afford to increase advertising at all. Indeed, firms cannot even afford to maintain their steady-state level of advertising, such that advertising spending falls by roughly 1.5% upon the initial drop in TFP. This means that output falls by (slightly) more in the initial period in this specification compared to the baseline case.

Moreover, the fall in advertising spending means that the stock of advertising falls (via depreciation) such that the return to having more products also falls. This leads to firms spending less on R&D so an insufficient number of products are introduced to cover the depreciation of the existing number of products, such that the aggregate stock of

products also falls. Hence, the returns to advertising are lower, so advertising spending also remains below the steady-state level for some time. Similarly, the lower advertising stock means aggregate R&D spending remains below the steady-state level for a long time, with the result that output and consumption take a long time to recover to the steady-state level. In fact, both only return to the steady-state level roughly 50 years after the initial drop in aggregate TFP. All of this is driven by the higher cost of advertising leading to firms not being able to ameliorate the fall in aggregate TFP in the first few years of the recession, and means that higher advertising effectiveness allows the economy to recover more quickly from a recession.

A different mechanism occurs when advertising effectiveness is decreased by reducing the exponent α on advertising in the CES aggregator. In this case, firms spend even more on advertising upon the initial drop in aggregate TFP (compared to the baseline case) and cut their R&D spending by even more in order to fund it (R&D and advertising are still substitutes at this point of the recovery, similar to the baseline case). This is because firms can afford to increase advertising but need to increase it by even more in order for it to ameliorate the drop in TFP. Indeed, they increase their advertising spending by such an amount that the initial drop in output is (slightly) less in this case compared to the baseline case. However, this means that the over-correction described for the baseline case is even more prevalent here.

As in the baseline case, a few years after the initial drop in aggregate TFP, R&D and advertising become complementary such that the substantially higher advertising stock means that R&D spending and the aggregate stock of products increase by even more compared to the baseline case, such that output exceeds the steady-state level even more quickly when advertising is less effective. Note, however, that consumption recovers more slowly when advertising is less effective, due to R&D and advertising spending being so much higher. The economy then returns to the steady-state level much more quickly: output and consumption are at their steady-state level roughly 30 years after the initial shock, compared to 50 years in the baseline case. As such, when advertising has a lower effect on prices and demand through the CES aggregator, the economy's response to a drop in aggregate TFP is much more volatile.

Finally, when both costs are increased and the exponent on advertising in the CES aggregator is reduced, firms have almost no ability nor incentive to try to counteract the drop in aggregate TFP by increasing their advertising spending. Hence, even though the initial drop in the aggregate variables is roughly similar as in the other scenarios for the first few periods, the path is very different after that. In particular, advertising spending drops substantially after period five, as does the stock of advertising. Both variables reach a trough between periods ten and fifteen. This means that firms' returns to developing new goods are also much lower, such that their investment in R&D and

the number of products falls substantially more than in the other scenarios. This means that output and consumption are far below the recovery paths in the other scenarios from around period five until period twenty.

However, because this scenario has a low stock of advertising in the steady-state itself, the gradual recovery of this stock to its steady-state level from period twenty onward does not hinder the recovery of R&D investment, the number of goods, output, or consumption (because advertising has less of an impact on the marginal return to a new product in this scenario). Indeed, these latter four variables recover quickly back to their steady-state level, such that the economy returns to its steady-state at roughly the same time as the scenario in which just the CES exponent on advertising was reduced.

Overall, therefore, increases in the effectiveness of advertising appear to have made the economy's response to recessions caused by drops in aggregate TFP quicker and less volatile.

6. Conclusion

The effectiveness of advertising has increased substantially over the past fifty years, with the advent of the internet and the increased feasibility of highly-targeted advertisements. To investigate the impact this increase in advertising effectiveness has had on the US economy, both to the long-run steady-state and to the short-run recovery from a recession, I develop a discrete-time dynamic stochastic general equilibrium model of heterogeneous firms. Multi-product intermediate goods firms can choose to invest in R&D in order to increase their chances of developing a new product next period and / or invest in advertising to boost the demand for their pre-existing products. There is aggregate uncertainty in the form of shocks to aggregate TFP.

I model the change in effectiveness of advertising in three distinct ways: 1) a change to the cost of advertising; 2) a change to the impact a certain amount of advertising has on the demand for a firm's products; and 3) a combination of those first two scenarios. I find that the impact of all three approaches on the steady-state is very similar: more effective advertising means that high productivity firms do more R&D (and low productivity firms do less R&D) when advertising is more effective such that output and consumption are increasing in advertising effectiveness.

However, when examining the economy's recovery from a recession, the three approaches yield different insights: when advertising is more effective, the recovery is quicker (but more volatile) compared to when advertising costs more and less volatile (albeit slightly slower) compared to when advertising has a lower effect on demand for a firm's products. The recession is even deeper in the combined scenario when both when the cost of advertising is higher and advertising has less of an impact on demand.

This suggests that there is a trade-off between speed and volatility of the recovery when it comes to advertising effectiveness and that there might be a "sweet spot" that finds a balance between the two: the exact nature of this sweet spot and how to achieve it is left open for further work. The potential question of whether to subsidise or tax R&D and / or advertising to improve outcomes further is also left open.

References

Afrouzi, Hassan, Andres Drenik, and Ryan Kim. 2023. "Concentration, Market Power, and Misallocation: The Role of Endogenous Customer Acquisition *."

Aghion, Philippe, and Peter Howitt. 1992. "A Model of Growth Through Creative Destruction." *Econometrica*, 60: 323–351.

Akcigit, Ufuk, and William R Kerr. 2018. "Growth through heterogeneous innovations." *Journal of Political Economy*, 126: 1374–1443.

Argente, David, Doireann Fitzgerald, Sara Moreira, and Anthony Priolo. 2021. "How Do Firms Build Market Share?"

Argente, David, Munseob Lee, and Sara Moreira. 2024. "The Life Cycle of Products: Evidence and Implications." *Journal of Political Economy*, 132: 337–390.

Askenazy, Philippe, Thomas Breda, and Delphine Irac. 2015. "Advertising and R&D: Theory and Evidence from France."

Ayres, Joao, and Gajendran Raveendranathan. 2023. "Firm entry and exit during recessions." *Review of Economic Dynamics*, 47: 47–66.

Baslandze, Salome, Jeremy Greenwood, Ricardo Marto, and Sara Moreira. 2023. "The expansion of varieties in the new age of advertising." *Review of Economic Dynamics*, 50: 171–210.

Bernard, Andrew B, Stephen J Redding, and Peter K Schott. 2010. "Multiple-product firms and product switching." *American Economic Review*, 100: 70–97.

Broda, Christian, and David E Weinstein. 2010. "Product creation and destruction: Evidence and price implications." *American Economic Review*, 100: 691–723.

Cavenaile, Laurent, and Pau Roldan-Blanco. 2021. "Advertising, Innovation, and Economic Growth." *American Economic Journal: Macroeconomics*, 13: 251–303.

Cavenaile, Laurent, Murat Alp Celik, Jesse Perla, and Pau Roldan-Blanco. 2023. "A Theory of Dynamic Product Awareness and Targeted Advertising."

Cavenaile, Laurent, Murat Alp Celik, Pau Roldan-Blanco, and Xu Tian. 2024. "Style over substance? Advertising, innovation, and endogenous market structure." *Journal of Monetary Economics*.

Clementi, Gian Luca, and Berardino Palazzo. 2016. "Entry, Exit, Firm Dynamics, and Aggregate Fluctuations." *American Economic Journal: Macroeconomics*, 8: 1–41.

Cociuba, Simona E, Edward C Prescott, and Alexander Ueberfeldt. 2018. "US hours at work." *Economics Letters*, 169: 87–90.

Crouzet, Nicolas, and Janice C Eberly. 2019. "Understanding Weak Capital Investment: The Role of Market Concentration and Intangibles."

Dehez, Pierre, and Alex Jacquemin. 1975. "A Note on Advertising Policy Under Uncertainty and Dynamic Conditions."

Dekle, Robert, Atsushi Kawakami, Nobuhiro Kiyotaki, and Tsutomu Miyagawa. 2021. "Product Dynamics and Aggregate Shocks: Evidence from Japanese Product and Firm Level Data."

Dinlersoz, Emin M., and Mehmet Yorukoglu. 2012. "Information and industry dynamics."

Dixit, Avinash, and Victor Norman. 1978. "Advertising and Welfare."

Erdem, Tulin, and Baohong Sun. 2002. "An Empirical Investigation of the Spillover Effects of Advertising and Sales Promotions in Umbrella Branding." *Journal of Marketing Research*.

Fogg-Meade, Emily. 1901. "The Place of Advertising in Modern Business."

Gijsenberg, M J, H J Van Heerde, M G Dekimpe, and J.-B E M Steenkamp. 2009. "Advertising and price effectiveness over the business cycle."

Goldfarb, Avi. 2014. "What is Different About Online Advertising?" *Review of Industrial Organization*, 44: 115–129.

Hall, Robert E. 2012. "The cyclical Response of Advertising Refutes Counter-Cyclical Profit Margins in Favor of Product-Market Frictions."

Hamano, Masashige, and Francesco Zanetti. 2017. "Endogenous product turnover and macroeconomic dynamics." *Review of Economic Dynamics*, 26: 263–279.

Haskel, Jonathan, and Stian Westlake. 2018. *Capitalism Without Capital.* Princeton University Press.

Ignaszak, Marek, and Peter Sedlacek. 2023. "Customer Acquisition, Business Dynamism and Aggregate Growth *."

Jo, In Hwan. 2021. "Firm Size and Business Cycles with Credit Shocks."

Kaiser, Ulrich. 2005. "A Microeconometric note on product innovation and product innovation advertising." *Economics of Innovation and New Technology*, 14: 573–582.

Khan, Aubhik, and Julia K Thomas. 2013. "Credit shocks and aggregate fluctuations in an economy with production heterogeneity." *Journal of Political Economy*, 121: 1055–1107.

Klein, Michael A., and Fuat Åener. 2023. "Product innovation, diffusion and endogenous growth." *Review of Economic Dynamics*, 48: 178–201.

Klette, Tor Jakob, and Samuel Kortum. 2004. "Innovating Firms and Aggregate Innovation." *Journal of Political Economy*, 112: 986–1018.

Loecker, Jan De, Jan Eeckhout, and Gabriel Unger. 2020. "The Rise of Market Power and the Macroeconomic Implications." *The Quarterly Journal of Economics*, 135.

Luttmer, Erzo G J. 2011. "On the mechanics of firm growth." *Review of Economic Studies*, 78: 1042–1068.

Mcgrattan, Ellen R, and Edward C Prescott. 2014. "A Reassessment of Real Business Cycle Theory." Federal Reserve Bank of Minneapolis.

Minniti, Antonio, and Francesco Turino. 2013. "Multi-product firms and business cycle dynamics." *European Economic Review*, 57: 75–97.

Molinari, Benedetto, and Francesco Turino. 2018. "Advertising and Aggregate Consumption: A Bayesian DSGE Assessment." *Economic Journal*, 128: 2106–2130.

OECD. 2022. "OECD R&D tax incentives database, 2021 edition."

Qi, Shi. 2019. "Advertising, industry innovation, and entry deterrence." *International Journal of Industrial Organization*, 65: 30–50.

Ridder, Maarten De. 2024. "Market Power and Innovation in the Intangible Economy." *American Economic Review*, 114: 199–251.

Sedlacek, Petr, and Vincent Sterk. 2017. "The growth potential of startups over the business cycle." *American Economic Review*, 107: 3182–3210.

Shah, Rohan. 2024. "Entering Recovery."

Srinivasan, Raji, Gary L Lilien, and Shrihari Sridhar. 2011. "Should Firms Spend More on Research and Development and Advertising During Recessions?" *Journal of Marketing*.

Steenkamp, Jan-Benedict E M, and Eric Fang. 2011. "The Impact of Economic Contractions on the Effectiveness of R&D and Advertising: Evidence from U.S. Companies Spanning Three Decades." *Marketing Science*, 30: 628–645.

Tellis, Gerard J., and Kethan Tellis. 2009. "Research on Advertising in a Recession." *Journal of Advertising Research*, 49: 304–327.

Wernerfelt, Birger. 1988. "Umbrella Branding as a Signal of New Product Quality: An Example of Signalling by Posting a Bond."

Appendix A. Firm-level complementarity or substitutability of advertising and R&D

It is possible for advertising and R&D to be either complements to, or substitutes for, each other. They can be substitutes if firms cannot afford to spend on both and therefore in order to increase spending on one they must reduce spending on the other; or they can be complements if firms increase (or decrease) spending on both at the same time to heighten (or dampen) the effect they have on each other (perhaps because higher advertising can also increase the returns to R&D through increasing the demand for a new product and vice versa).

The possibility that the complementarity / substitutability of advertising and R&D differs across firms can be demonstrated using firm-level data from Compustat. I use Compustat North America annual data from 1980 until 2024. I remove all financial and regulated firms (i.e. all firms with SIC codes between 6000 and 6999, and between 4900 and 4999) and I exclude all observations that have missing or non-positive values for sales; total assets; number of employees; gross property, plant, and equipment; depreciation; operating income; capital expenditures; R&D expenditures; and advertising expenditures. I deflate all series using the US Producer Price Index.

I construct indicator variables as to whether a firm increased, decreased, or left unchanged their R&D expenditures compared to the previous year. I do the same for a firm's advertising expenditures. This resulted in each firm-year observation falling into one of nine different "buckets", which then provides some information as to whether or not a firm considered R&D and advertising as substitutes or complements (in that particular year), as shown in Table A.6. For example, if a firm changed its R&D and advertising expenditures in the same way from year-yo-year (i.e. increased both or decreased both at the same time, consistent with them being complements), but if it changed them in opposite directions (i.e. increased one while decreasing the other, consistent with them being substitutes for each other). Note that when a firm changes one of R&D and advertising but leaves the other unchanged, then we cannot infer as easily whether the firm considers them to be complements or substitutes.

Table A.7 shows the proportion of firm-year observations that are in each portion. These results suggest that at least 23% of firm-years (the 19.2% of firm-years that increased both plus the 3.9% of firm-years that decreased both) considered R&D and advertising to be complements, while at least 6.5% of firm-years (the 2.9% that increased advertising while decreasing R&D plus the 3.9% that decreased advertising and increased R&D) considered them to be substitutes. The remaining firms did not change at least one of the two.

This is strongly indicative of heterogeneity across firms in terms of whether R&D and advertising are complements or substitutes: in some cases, firms behave as though they are complements (changing both in the same direction) while in other cases they act

Table A.6: Implications of firm-level changes in R&D and advertising spending for complementarity / substitutability

	Increased R&D	Unchanged R&D	Decreased R&D
Increased Advertising	Complements	Unclear	Substitutes
Unchanged Advertising	Unclear	Unclear	Unclear
Decreased Advertising	Substitutes	Unclear	Complements

Table A.7: Proportion of firm-year observations with respective year-to-year changes in R&D and advertising

	Decreased R&D	Unchanged R&D	Increased R&D
Increased Advertising	2.6	9.0	19.2
Unchanged Advertising	10.3	29.4	14.8
Decreased Advertising	3.9	6.8	3.9

as though they are substitutes (increasing one while reducing the other). As such, it is crucial to allow for this heterogeneity in any model of R&D and advertising.

Appendix B. Solving the model

In this Appendix I first set out the solution algorithm I use to solve for the non-stochastic steady-state before describing the solution method I use to obtain my results with aggregate uncertainty in the form of shocks to aggregate TFP. Throughout, a firm's optimum labour-hiring is given by the solution to $\pi(a, \varepsilon_o, z, \mu) = \max_{\ell,p} p(a, \varepsilon_o, z, \mu) y(z, \mu) - w(z, \mu)\ell$ where $y = z\varepsilon_o \ell^{\gamma}$. This results in optimum labour hiring given by

$$l^*(a, \varepsilon_o, z, \mu) = \left[\frac{w(z, \mu)}{\gamma \frac{\theta-1}{\theta} (1 + \frac{a}{A})^{\alpha} Y(z, \mu)^{\frac{1}{\theta}} (\varepsilon_o z)^{\frac{\theta-1}{\theta}}} \right]^{\frac{\theta}{\gamma\theta-\gamma-\theta}}$$

I then plug this back into Equation (7) when solving for firms' choices of advertising a' and R&D R.

Appendix B.1. Steady-state

Individual firms' decisions of how much to produce depend on the wage (which itself depends on aggregate consumption as discussed in Section 2.1) and aggregate output. In addition, the amount of advertising with which entrants start is some fraction ψ of aggregate advertising. Therefore, the first step is to make guesses of aggregate output, consumption, and aggregate advertising.

Next, for each starting combination of the number of products and output productivity, I obtain a firm's optimum choice of advertising expenditure for the next period I_a conditional on its choice of investment (assuming that the firm can afford both without paying negative dividends; i.e. that it is "unconstrained"). A firm's choice of advertising expenditure for the next period is given by maximising Equation (7) with respect to a' and then applying the Benveniste-Shenkman result to obtain an expression for the derivative of the value function in the next period. This gives the optimum stock of advertising for the next period, conditional on the choice of R&D investment (and the current firm's own and aggregate stock of advertising) as:

$$a' = A'\left[\left(\frac{\frac{q(z,\mu)}{\beta\varepsilon_a}\cdot\frac{\gamma\theta-\gamma-\theta}{-\alpha\theta}}{[\Phi(R)\frac{((1-\delta_n)N+1)}{A'}\sum q(z',\mu')\Omega+(1-\Phi(R))\frac{(1-\delta_n)N}{A'}\sum q(z',\mu')\Omega]}\right)^{\frac{\gamma\theta-\gamma-\theta}{-\alpha\theta-\gamma\theta+\gamma+\theta}}-1\right]$$

where

$$\Omega = Y^{\frac{1}{\theta}}(\varepsilon_o z)^{\frac{\theta-1}{\theta}}\left[\frac{w}{\gamma(\frac{\theta-1}{\theta})(1+\frac{a}{A})^{\alpha}Y^{\frac{1}{\theta}}(\varepsilon_o z)^{\frac{\theta-1}{\theta}}}\right]^{\frac{\gamma(\theta-1)}{\gamma\theta-\gamma-\theta}} - w\left[\frac{w}{\gamma(\frac{\theta-1}{\theta})(1+\frac{a}{A})^{\alpha}Y^{\frac{1}{\theta}}(\varepsilon_o z)^{\frac{\theta-1}{\theta}}}\right]^{\frac{\theta}{\gamma\theta-\gamma-\theta}}$$

Then, the unconstrained firm's optimum choice of advertising expenditure is simply determined as $I_a = \max(a' - (1-\delta_a)a, 0)$. Note that as the optimum unconstrained choice of a' depends on the probability of a successful innovation $\Phi(R)$, which itself is a function of R&D investment, the unconstrained choice of I_a depends on the choice of R&D investment.

In the next step, I use this expression to solve for the decision rules via Value Function Iteration (VFI) for just "unconstrained firms using a single variable Golden Section Search over R&D investment R. In other words, I assume that all firms can afford the optimum choice of R&D and advertising expenditure and search over their choice of R&D investment (obtaining the implied choice of I_a via the method above) and obtaining the value of that choice.

Next, I identify which firms are unconstrained and those which are "constrained": the former are able to afford the unconstrained optimum choices of R&D and advertising expenditure found in the previous step without violating the non-negative dividend constraint. Constrained firms, on the other hand, are not able to afford these choices without paying negative dividends. These constrained firms are then further split into two types: those that can afford the fixed cost of R&D ξ_R and those that cannot.

For the firms that can afford the fixed cost of R&D, their decision over the two choice variables R&D investment and advertising expenditure for the next period can be reduced to a Golden Section Search over just R&D investment with their advertising expenditure then given by $I_a = \varepsilon_a(N\pi(a, \varepsilon_o, z, \mu)_i - \frac{(1-\tau_r)R}{\varepsilon_r} - \mathbb{I}_r\xi_r$. For the firms that cannot afford the fixed cost of R&D, their decision over the same two choice variables is obtained by fixing their R&D investment equal to zero and just conducting a Golden Section Search over their choice of advertising expenditure. Using this, I then conduct a sequential value function iterations (using the value function found in the previous step as the starting point) and find the decision rules for all constrained firms (leaving those for unconstrained firms unchanged) to get a value function for all firms.

Using these decision rules I iterate on the distribution until it converges and use that distribution along with firms' decision rules to obtain a new value for aggregate output and consumption. The initial guesses of these variables are then updating using Broyden's algorithm until the initial guess and the new values converge.

Appendix B.2. With aggregate uncertainty

The presence of aggregate uncertainty in the form of varying aggregate productivity z (which can take three values) means that I use the Krusell-Smith solution method to solve my model, adapted for heterogeneous firms as per Khan and Thomas (2013). This means that I approximate the distribution of firms μ using the summary statistic of the total number of products $\bar{N}$ and its lag $\bar{N}_{t-1}$, the aggregate advertising stock A, the lag of the value of the exogenous TFP shock z_{t-1}, and the lag of the number of products held by single-good firms $g_{t-1}$and use that in the forecasting rules that take the form

$$\log x = \beta_{0,i} + \beta_{1,i} \log \bar{N} + \beta_{2,i} \log A + \beta_{3,i} \log \bar{N}_{t-1} + \beta_{4,i} \log z_{t-1} + \beta_{5,i} \log g_{t-1} + v_i$$

where x is the variable to be forecast, β_0, β_1, β_2, β_3, $\beta_4$and β_5 are coefficients to be estimated, v is the error term, and i represents the level of aggregate productivity in the current period. Each forecast rule is obtained via regressions using the final 1,000 observations from a simulation of 1,500 periods (i.e. the first 500 periods of the simulation are dropped when estimating the forecast rules).

As mentioned in the main text, I require forecasting rules for aggregate output Y (as that affects the price intermediate firms charge for their output, thereby affecting the output of those intermediate firms) and the household's marginal utility of consumption q (to determine the marginal utility value of firm profits today, as firms are owned by the household) in the current period. I also need forecasting rules for the number of products next period $\bar{N}'$ and the aggregate stock of advertising next period A'.

The inner loop of my solution algorithm uses a guess of the forecasting rules for each of these four variables alongside the conditional optimum advertising expenditure and value function iteration procedures described in Appendix B.1. In an outer loop, I simulate the economy for 1,500 periods, using Broyden's algorithm each period to obtain the marginal utility of consumption q and aggregate output Y that prevails in that period (and which firms use in their decision making). I then estimate new forecast rules dropping the first 500 observations of the simulated economy and update the guess of the forecasting rules. This iterative procedure continues until convergence.

The estimated coefficients, and some regression statistics for each of these forecast rules for the full model are shown in Table B.8. The first column indicates which variable is being forecast and the second column shows to which level of the aggregate productivity in the current period the forecast rule applies. Columns three to eight contain the estimated coefficient values. The final two columns contain statistics that represent the accuracy of these forecast rules: column nine shows the standard error of the regression itself and the final column shows the R^2 of the forecast rules.

In each forecasting rule, and for each value of the aggregate TFP shock z, the R^2 is almost always above 0.99 (the only exception is the forecasting rule for the future stock of adverts A' when current exogenous TFP is at its median level and even in that instance the $R^2$is relatively close to 0.99). In other words, these forecasting rules are accurate in terms of agents in the model correctly forecasting the simulated economy's actual state. Although I only present forecasting rules for the baseline model here I obtain new forecast rules for each variation of the effectiveness of advertising by repeating the solution algorithm described above for each specification (not presented here) and each of those obtains similar levels of accuracy as the forecasting rules for the baseline model.

Appendix C. Model validation against untargeted moments

As one method of externally validating of my model I examine the extent to which it can capture the firm-size distribution, which is shown in Figure C.9. In this graph, the solid black bars represent the distribution of firms over bins defined using the number of people they employ obtained from BDS data averaged over the period 1980-2006. The grey patterned bars show the distribution of firms over these bins that result from my

Table B.8: Forecast rules for the full model

Forecast rule for	Aggregate productivity	β_0	β_1	β_2	β_3	β_4	β_5	SE	R^2
$\bar{N}'$	z_1	0.293	1.359	0.057	-0.650	-0.004	-0.148	4.73E-04	1.000
	z_2	0.393	1.333	0.102	-0.639	-0.011	-0.213	5.51E-04	0.999
	z_3	0.369	1.311	0.089	-0.609	-0.008	-0.186	6.07E-04	0.999
A'	z_1	-1.331	0.584	0.411	-0.332	0.088	-0.394	6.29E-04	0.993
	z_2	-1.321	0.917	0.435	-0.539	0.104	-0.497	2.29E-03	0.967
	z_3	-1.409	0.478	0.371	-0.205	0.071	-0.391	6.80E-04	0.992
q	z_1	0.722	-0.003	0.031	-0.334	0.022	-0.161	2.83E-04	0.999
	z_2	0.693	-0.041	0.029	-0.303	0.017	-0.163	2.96E-04	0.999
	z_3	0.653	-0.036	0.021	-0.306	0.016	-0.153	2.65E-04	1.000
Y	z_1	-0.564	0.377	-0.019	0.008	-0.012	0.009	2.53E-04	1.000
	z_2	-0.546	0.376	-0.020	0.014	-0.010	0.007	2.34E-04	1.000
	z_3	-0.522	0.382	-0.020	0.013	-0.007	0.009	1.95E-04	1.000

Note: SE is the standard error in each regression

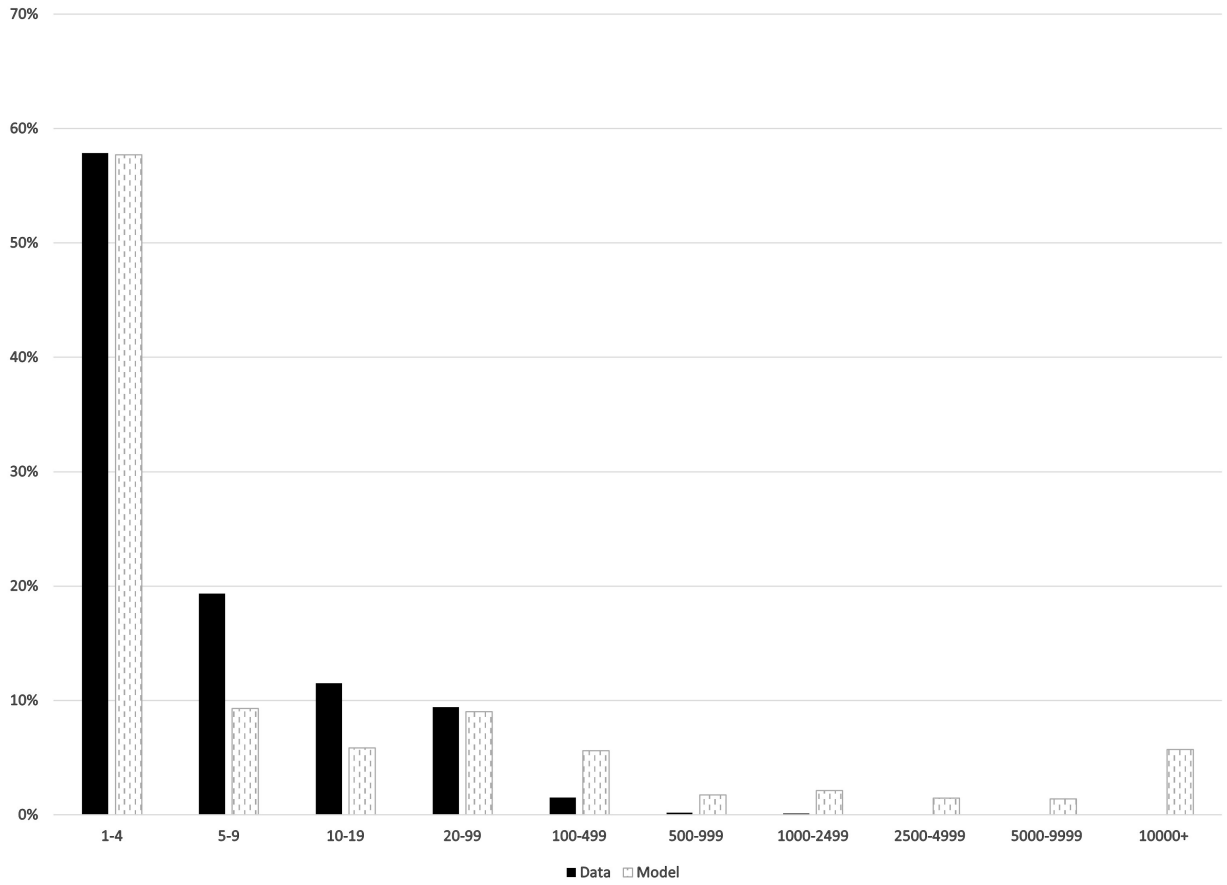
model, with those bins defined using the same approach as Jo (2021).[31] My model is able to get very close to the Pareto shape of this (untargeted) distribution despite using only (log-) normally distributed idiosyncratic productivity shocks. In particular, it exactly matches the proportion of firms in the smallest employment bin but has a few too many large firms (those with more than one hundred employees) at the expense of medium-sized ones (those with between five and nineteen employees).

My model also captures the fact that entrants grow much more quickly than do old firms, as found by Luttmer (2011). In particular, Figure C.10 plots average variables for a cohort of entrants over time after their entry. For each variable, the y axis shows the value of the variable relative to the value it took in the cohort's year of entry. The top left panel plots the cohort's average sales (defined as $Np(a, \varepsilon_o, z, \mu)y(\varepsilon_o, z, \mu)$) while the bottom left panel shows the cohort's average employment. The middle column of panels show, respectively, the cohort's stock of goods and their total R&D spending, while the rightmost column shows the cohort's stock of advertising and their advertising expenditure.

In each variable, conditional on surviving (i.e. not being hit with the exogenous death shock) the average new entrant grows quickly until it is roughly 20 years old. At this

[31] Jo (2021) first uses the cumulative proportion of employment in each bin to convert the "real life" employment range of each bin to something that can be used within the model. Then the proportion of firms that falls within each of these model-defined bins is calculated to obtain the overall firm-size distribution.

Figure C.9: Comparison of model to untargeted size-distribution of firms



Note: The solid black bars show the distribution of firms over employment obtained using BDS data for the period 1980-2006. The grey patterned bars show the distribution of firms over employment for the non-stochastic steady-state of my model, with the proportion of firms in each bin obtained following the procedure described in Jo (2021) to determine the proportion of employment accounted for each bin, using the cumulative distribution of employment to define each bin within the model, and then calculating the proportion of firms that are within each such-defined bin.

point its growth slows down and firms slightly decrease in size over the next 200 years. In other words, smaller, younger firms grow quickly relative to larger, older firms, exactly capturing the path of firm growth shown in Luttmer (2011) and Cavenaile and Roldan-Blanco (2021). Moreover, the rapid increase in the average (surviving) entrants' sales in the first few years of its existence captures the path of sales in the first 3-4 years as found by Argente et al. (2021). As such, my assumption that firms can only add one product (at most) in any one year ensures that my model capturing the rapid growth of small, young firms and the relatively slower growth of larger, older firms.

Focusing on the cohort's stock of advertising and their advertising expenditure, the hump-shaped nature of these variables as the cohort ages (i.e. a peak in the stock of advertising and advertising expenditure around 20 years before a gradual decline as the cohort ages) matches Argente et al. (2021)'s result that a firm's advertising spending increases in the first few years of a firm's entry before decreasing after that. Moreover, it also matches Klein and Åener (2023)'s finding that new firms advertise substantially to promote their new existence and raise customer awareness before reducing their

Figure C.10: Growth of a cohort of entrants



Note: Each line shows the average variable per firm for a simulated cohort of entrants where the level of consumption and output remains at the steady-state level throughout the simulation (note that a firm's sales is given by its price times output times number of products). Each line is relative to the value it takes during the year of the cohort's entry.

advertising expenditures as they age.

Finally, my model also gets close to capturing Cavenaile and Roldan-Blanco (2021)'s result that the ratio of a firm's advertising expenditure to its sales decreases as the firm's size increases. The comparison for this ratio is shown in Table C.9, with the moments for the data obtained from Compustat.[32] I express each figure relative to the first quintile to enable comparison between the model and data. The data exhibit a large drop in the advertising / sales ratio between the first and second quintile and then has a slight "u-shape" for the top quintiles as the ratio for the fifth quintile is slightly higher than that for the fourth quintile. Despite not targeting this moment or pattern, my model captures the large drop between the first and second quintile as well as the slight increase between the fourth and fifth quintiles, only missing off the data being relatively flat between quintiles

[32] The data moments are obtained from Compustat annual data since 1980 having dropped all financial and regulated firms from the sample. Outlier observations in which a firms sales, assets, employees, "property, plant, and equipment", depreciation, accumulated depreciation, or capital expenditures are recorded as being negative are also dropped from the sample. I then allocate each firm in each year to a quintile based on its sales in that particular year relative to all other firms in the sample for that year. I then calculate each firm's advertising expenditure divided by its sales and drop all observations where this ratio is greater than 1 before finding the average ratio for each quintile across all years in the sample.

two and three (with my model having another relatively large decrease between those quintiles).

Table C.9: Untargeted moment: advertising expenditure / sales

Sales Quintile	Model	Data
1	1	1
2	0.501	0.692
3	0.389	0.689
4	0.613	0.609
5	0.614	0.616

Note: The data moments are obtained from Compustat annual data since 1980 having dropped all financial and regulated firms from the sample. Outlier observations in which a firms sales, assets, employees, "property, plant, and equipment", depreciation, accumulated depreciation, or capital expenditures are recorded as being negative are also dropped from the sample. I then allocate each firm in each to a quintile based on its sales in that particular year relative to all other firms in the sample for that year. I then calculate each firm's advertising expenditure divided by its sales and drop all observations where this ratio is greater than 1 before finding the average ratio for each quintile across all years in the sample.

Nonetheless, my model is able to capture each of these aspects of the data relatively well despite not targeting them.

Appendix D. Validation of model business cycle statistics

As described in Section 3.2, I simulate the full model for 1,500 periods and then examine the economy from the five-hundredth period onward results in the business cycle statistics. These results are presented in Table D.10. The first row shows the mean of each series, while the second shows their coefficient of variation. The third row shows the volatility of each series relative to that of output. The fourth row shows each series' correlation with output, while the bottom row shows each series' correlation with its own one-year lag.

The respective values of the bottom three rows (i.e. the volatility of each series relative to that of output, each series' correlation with output, and each series' auto-correlation) for the US economy itself are shown in Figure D.11 (note that no data were available for the total number of products in the US economy). Each of these results is obtained by analysing the cyclical component obtained from applying an Hodrick-Prescott filter with a smoothing parameter of 400 to each series. Note that it is not possible to compare the mean and coefficient of variation between the model and the US economy as a whole precisely because the levels of each series obtained in the model are substantially different from those seen in the data (for example, output in the model is not in the billions). However, it is possible to compare the other three metrics as these account for the fact that the level of each series differs between the model and the US economy as a whole.

Table D.10: Simulated business cycle moments for baseline model

	Y	C	$\bar{N}$	Total R&D Spending	A	Total Advertising Spending
Mean	0.79	0.68	1.89	0.014	0.12	0.024
Coefficient of Variation	2.7%	2.7%	2.8%	4.0%	1.5%	3.3%
$\frac{\sigma_X}{\sigma_Y}$	(0.021)	0.87	2.52	0.03	0.08	0.04
Correlation with GDP	1.00	1.00	0.76	0.94	0.77	0.55
Auto-correlation	0.88	0.88	0.98	0.96	0.89	0.09

Note: These moments are obtained from a simulation of 1,500 periods of the full model, dropping the first 500 periods of that simulation. These moments have not been filtered or smoothed. The first two rows show, respectively, the mean and coefficient of variation of each series. The third row shows the standard deviation of each series relative to that of output, except for the first column, which shows the actual standard deviation of output. The fourth row shows each series' contemporaneous correlation with output, while the last row contains each series 1-period auto-correlation. All numbers are expressed in levels.

Table D.11: Business cycle moments for the US economy

	Y	C	$\bar{N}$	Total R&D Spending	A	Total Advertising Spending
$\frac{\sigma_X}{\sigma_Y}$	(240.34)	0.58	N/A	0.05	0.07	0.06
Correlation with GDP	1.00	0.94	N/A	0.38	0.85	0.83
Auto-correlation	0.67	0.62	N/A	0.80	0.78	0.69

Note: These moments are obtained by analysing the cyclical component obtained from applying a Hodrick-Prescott filter with a smoothing parameter of 400 to each series (which is recorded annually). The first row shows the standard deviation of each series relative to that of output. The second row shows each series' contemporaneous correlation with output, while the last row contains each series 1-period auto-correlation. All numbers are expressed in levels. The source for output is the constant 2017 USD output series from the NSF's National patterns dataset, while that for consumption is FRED's household consumption expenditures minus FRED's household durable goods spending, deflated using the gdp deflator. The series on private R&D spending is total private R&D spending from the NSF, deflated using the GDP deflator. The advertising spending series is created by taking total US advertising spending over the period 1900-2007 from https://www.galbithink.org/ad-spending.htm and then extending the series until 2020 using IRS data on reported corporate income tax returns. The two series overlap for three years (between 2005 and 2007) and are consistent with each other during that overlap period. The advertising stock series is obtained by applying the perpetual inventory method to the advertising spending series, assuming an annual depreciation rate of 75% and an starting value for the stock of the advertising spending in that initial year divided by the depreciation rate.

None of these moments were targeted as part of my model's calibration. Nonetheless, my model performs well in matching the volatilities of consumption, R&D spending, advertising spending, and the stock of advertising relative to output (the first row). Indeed, my model performs well in capturing the co-movements of each series with output as well as the auto-correlation of each series (the only slight "misses" of my model are that R&D spending in my model is more correlated with output, and advertising spending is not sufficiently correlated with itself, compared to the data). As such, my model performs well at capturing the path of the US over the business cycle despite not explicitly targeting any of these moments.

Appendix E. Firms switch from spending on R&D to spending on advertising at the start of a recession

In my baseline model, firms switch from investing in R&D to spending more on advertising in the first year of a recession. This is consistent with the data as shown below in Table E.12. The first column of this table shows the first year of an NBER-identified recession, while the second column shows the percentage point change in the cyclical component of aggregate advertising spending in that year compared to the previous year. The third column shows the corresponding percentage point change in the cyclical component of aggregate R&D spending. In each case, the cyclical component has been obtained using a linear trend. The last row of the table shows the mean across all NBER recessions.

Table E.12: Percentage point changes in cyclical components of aggregate advertising and R&D in the first year of recessions

NBER Recession	% point change in cyclical component of advertising	% point change in cyclical component of R&D
1957	-6.8	-8.3
1960	-2.6	-6.0
1969	-2.4	-4.0
1973	-2.2	-3.5
1980	-2.8	-3.5
1990	-1.9	-2.3
2001	-8.9	-9.2
2007	-4.5	-4.9
2020	-15.7	-16.7
Mean	-5.3	-6.5

Note: The table shows the percentage point change in the cyclical components of aggregate advertising spending and aggregate R&D spending in the first year of NBER-identified recessions. The advertising data are created by taking total US advertising spending over the period 1900-2007 from https://www.galbithink.org/ad-spending.htm and then extending the series until 2020 using IRS data on reported corporate income tax returns. The two series overlap for three years (between 2005 and 2007) and are consistent with each other during that overlap period. The R&D series is obtained from the NSF's total business spending on R&D. Both series are detrended using a linear trend.

The first row of the table refers to the first year of the 1957-1958 NBER recession - i.e. it refers to 1957. It shows that the cyclical component of advertising decreased by 6.8% points in 1957 compared to 1956. In other words, advertising spending was 6.8% points more below its trend in 1957 than it was in 1956. Similarly, R&D spending was 8.3% points more below its trend in 1957 than it was in 1956. This is consistent with firms switching away from R&D towards advertising at the start of the recession: R&D spending fell by more relative to its trend than advertising spending did relative to its

trend, suggesting that firms tended to place a higher priority on cutting R&D spending and trying not to cut (as much of) their advertising spending at the onset of the 1957-1958 recession.

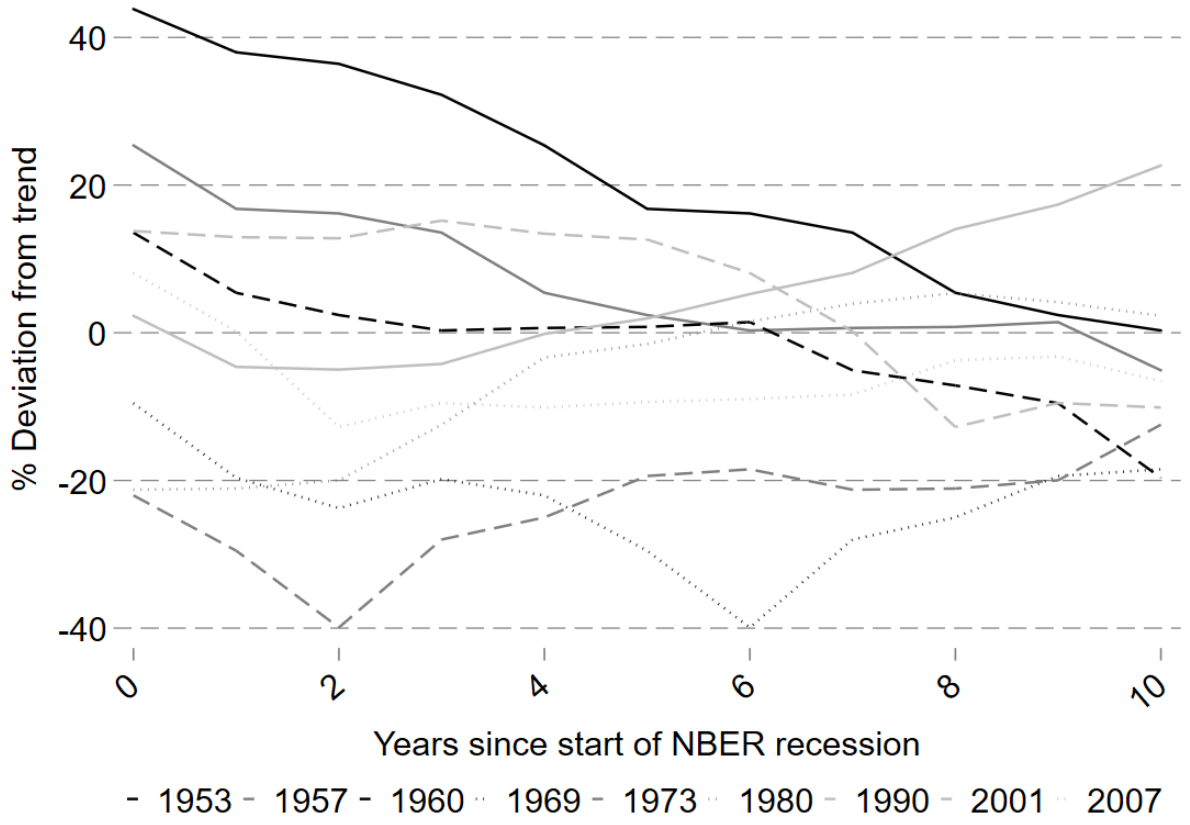
This is also the case for all of the recessions presented in the table: aggregate R&D spending decreased more relative to its trend than did aggregate advertising relative to its own trend in the first year of each recession. This is also true when calculating the mean across the recessions. In other words, at the onset of every recession, firms switched away from spending on R&D in favour of spending on advertising, matching the aggregate response to a TFP shock as found in my baseline model.

Appendix F. Advertising Spending is above trend in the early periods of recessions

As mentioned in Section 5.1 advertising expenditure tends to be above trend its trend in the early period of a recession. This can be seen in Figure F.11, which plots the cyclical component of advertising in the years after an NBER recession. Specifically, the y axis shows the percentage deviation from a linear trend while the x axis shows the number of years since the start of an NBER recession. Each line in the graph represents the cyclical path of advertising from the start of the respective NBER recession.

The majority of the series are above the zero-line in the initial years of the recession and recovery: five out of the nine recessions are above trend in the first year after the start of the recession, while at least four of the nine recessions are above trend in subsequent years. This indicates that the response of advertising spending in my baseline model (being above trend during the economy's recovery) matches at least some previous recessions in the US.

Figure F.11: Cyclical component of advertising spending in the years after the start of NBER recessions



Note: The graph shows the cyclical component of advertising in the years after an NBER recession. The y axis shows the percentage deviation from a linear trend while the x axis shows the number of years since the start of an NBER recession. Each line in the graph represents the cyclical path of advertising from the start of the respective NBER recession. The advertising series is created by taking total US advertising spending over the period 1900-2007 from https://www.galbithink.org/ad-spending.htm and then extending the series until 2020 using IRS data on reported corporate income tax returns. The two series overlap for three years (between 2005 and 2007) and are consistent with each other during that overlap period. The advertising series is then detrended using a linear trend.