Women Inventors and Economic Growth

Jean-Felix Brouillette*

March 12, 2024

Abstract

In 1976, 4% of inventors in the U.S. were women, and by 2020, that fraction had only moved up to 12%. Under the natural assumption that there are no intrinsic differences in inventive potential across genders, the scarcity of women in innovation reveals that the U.S. is missing out on some of its brightest minds. This raises two questions: (1) What are the barriers faced by those "lost" Jennifer Doudnas? and (2) How costly is the resulting misallocation of inventive talent for aggregate productivity and welfare? To tackle those questions, I develop a theory of semi-endogenous growth in which individuals with heterogeneous talent choose between a career in research or production. However, three gendered barriers can deter or prevent women from pursuing their comparative advantage. They may face different forms of discrimination in the labor market, be confronted with higher obstacles to human capital formation or lack the opportunities and role models to become innovators. Interpreting micro-level data on the universe of U.S. inventors through the lens of this framework, I find that the underrepresentation of women in innovation is virtually all due to a lack of exposure to innovation. Women and men inventors are just too similarly productive and educated for distortions operating through selection or human capital to play a prominent role. Taking advantage of the structure of this theory, I find that lifting all barriers to female innovation would increase U.S. income per person by 8.6% in the long run. Accounting for transition dynamics reveals that this policy would be equivalent to permanently raising everyone's consumption by 2.7%.

*Stanford University, Department of Economics, E-mail: jfbrou@stanford.edu. I am grateful to Pete Klenow, Chad Jones, Chris Tonetti, Monika Piazzesi, Martin Schneider, Adrien Auclert, Ernest Liu, Shifrah Aron-Dine and Mohamad Adhami for their helpful comments. This project was supported by the B.F. Haley and E.S. Shaw Fellowship for Economics at SIEPR, Stanford and the "Unlocking Our Potential" dissertation program at the Federal Reserve Board of San Francisco.

1 Introduction

Economic growth is the product of a handful of people discovering infinitely usable ideas that raise everyone's living standards. With this crucial role, the hope is that our brightest and most creative minds engage in the process of innovation. But is that really the case? In 1976, 4% of inventors in the U.S. were women, and by 2020, that fraction had only inched up to 12%.[1] Under the natural assumption that there are no *innate* gender differences in inventive ability, the vast underrepresentation of women in research reveals that the U.S. is missing out on some of its most talented inventors. This observation prompts two questions: (1) What are the barriers faced by those "lost" Jennifer Doudnas? and (2) How costly is the resulting misallocation of inventive talent for aggregate productivity and welfare?

To tackle those questions, I interpret micro-level data on the universe of U.S.-based inventors through the lens of a general equilibrium theory of semi-endogenous growth with expanding varieties. In this model, there are two occupations: inventors who discover new varieties and workers who produce existing ones. Individuals in this economy are heterogeneous in their raw talent for innovation and must make two key decisions: choose how much human capital to accumulate and decide whether to pursue a career in research or production (Roy, 1951). Yet, three occupation-specific barriers may deter or prevent women from pursuing their comparative advantage.

First, women can be denied due compensation for their inventions, which is modeled as an exogenous tax on their earnings. To illustrate this distortion, consider the case of Gerty Cori, the first American woman to win a Nobel Prize in science. Two years after discovering the Nobel Prize-winning Cori cycle with her husband Carl Cori, the only employment she could find was that of a research associate at Washington University, receiving a tenth of her husband's salary. While Carl was extended an invitation to chair the university's pharmacology department, Gerty would have to wait another 16 years to be promoted to a full professorship (American Chemical Society, 2004a,b).[2] Her unfortunate experiences echo the many forms of workplace discrimination that women confront throughout their careers in research. Some revealing examples are the asymmetrical sharing of rents from patents (Kline, Petkova, Williams and Zidar, 2019), unequal access to patent assistance or financial resources (Ewens and Townsend, 2020; Hannon, 2021; Pairolero, Toole, DeGrazia, Teodorescu and Pappas, 2022; Morazzoni

[1] Inventors are defined as residents of the United States who have been granted a patent by the U.S. Patent and Trademark Office. This study focuses on gender because other demographic characteristics such as race or family income are not available in patent data and are much harder to infer.

[2] Despite all of their work being collaborative, Carl Cori was also the sole recipient of the Albert Lasker Award for Basic Medical Research and the American Chemical Society's Willard Gibbs Medal.

and Sy, 2022) or outright denial of their scientific contributions (Jensen, Kovács and Sorenson, 2018; Hofstra, Kulkarni, Galvez, He, Jurafsky and McFarland, 2020; Hochberg, Kakhbod, Li and Sachdeva, 2023; Ross, Glennon, Murciano-Goroff, Berkes, Weinberg and Lane, 2022).

Second, women can face higher barriers to human capital formation throughout the entire innovation pipeline. This is modeled as differential disutility from time spent in schooling for aspiring women inventors. This distortion should be interpreted broadly, as a reflection of the many forms of discrimination that women and girls may face in the development of particular skills required in research careers. For instance, gender norms and stereotypes might influence how educators and parents invest in their students' or daughters' human capital by distorting both the extent and direction of those investments. Women may also be dissuaded from pursuing studies in fields where the learning environment is particularly hostile towards them. A recent report by the National Academies of Sciences, Engineering, and Medicine (2018) found that 20% to 50% of female students in U.S. STEM programs experienced sexual harassment from faculty or staff depending on their field and degree.

Lastly, women may simply not be as frequently "exposed" to inventive careers and opportunities during their formative years as their male counterparts, *regardless* of their talent. Receiving exposure to innovation is modeled as a binary random variable that determines whether someone can choose to invent or not. This distortion is intended to capture one of the key findings from Bell, Chetty, Jaravel, Petkova and Van Reenen (2018), which is that girls who grow up surrounded by more women who patent in a specific field are more likely to go on to patent in the same field. Hence, a lack of relevant role models could prevent girls from forming aspirations and/or having enough information to invent. Gerty Cori was no exception to this pattern. Her father Otto Radnitz was a chemist turned manager of beet-sugar refineries after successfully devising a sugar refining process. Perhaps not surprisingly, Gerty Cori would years later win her Nobel Prize for describing the fate of sugar in the human body.

Discerning the salience of each distortion is crucial for devising practical solutions to the gender gap in innovation. Not all barriers are created equal, and their respective toll on aggregate productivity necessitates a nuanced assessment of their consequences. Labor market discrimination and barriers to human capital accumulation mainly operate through selection, thereby deterring women from the lower echelons of inventive talent. Hence, were these distortions to be neutralized, we may not witness an unprecedented influx of highly talented women into the innovation process. In contrast, insufficient exposure to inventive careers indiscriminately closes the door to invention for women, irrespective of their position on the talent spectrum. Mitigating this particular barrier could thus unlock a vast reservoir of untapped talent, including potential innovators of exceptional ability. The implications are profound: targeting the right barriers could not only balance gender representation but also elevate the caliber of inventions and, consequently, propel economic growth.

To disentangle those distortions through the prism of the model, I rely on two key assumptions: (1) women and men draw their talent from the same distribution and (2) the probability of receiving exposure to innovation is independent of one's talent draw. While the first assumption is a natural starting point, the second is motivated by evidence from Bell et al. (2018) that children who grow up in areas with a greater concentration of inventors do not seem to have different latent cognitive abilities as proxied by early childhood test scores. With these premises, the model has clear implications for both the extensive (quantity) and intensive (quality) margins of innovation by gender, as well as women inventors' educational attainment, relative to their male colleagues. These three distinct moments provide insights into the salience of each distortion.

On the extensive margin, all barriers considered above shrink the pool of potential women inventors. Yet, only the labor market tax distorts the intensive margin of research labor supply. Indeed, by deterring women at the lower spectrum of inventive talent, this distortion increases female research productivity through positive selection. While the human capital barrier shares this selective mechanism, it also directly dampens research productivity by discouraging the accumulation of skills required in inventive careers. These counteracting *selection* and *direct* effects exactly offset one another to leave female research productivity unchanged. Similarly, since exposure to innovation is random and uncorrelated with research potential, it has no compositional effect on the pool of inventive talent. In that sense, the human capital and exposure distortions are observationally equivalent with respect to both margins of female research labor supply. However, the key distinction lies in the fact that the human capital distortion explicitly deters education, suggesting that women inventors would tend to achieve lower educational attainment on average.

Therefore, with sensible empirical counterparts to both the extensive and intensive margins of research effort as well as innovators' educational attainment, one can hope to shed light on the barriers faced by women inventors and discipline the model's key driving forces. The data on those first two moments comes from PatentsView, which contains information on all patents granted by the U.S. Patent & Trademark Office (USPTO) since 1976. PatentsView uses a series of disambiguation algorithms to uniquely identify inventors over time and, most importantly, predict their gender from their first name. Restricting on U.S.-based inventors delivers a large sample of 1.7M inventors to whom 3.7M patents have been granted. For insights into researchers' educational trajectories, I turn to the National Surveys of College Graduates (NSCG), narrowing the sample to respondents primarily engaged in research and development (R&D).

Analysis of the data reveals a pronounced underrepresentation of women among U.S. patent holders. However, when measuring inventive productivity–defined as the number of patents granted within a year, weighted by their stock market valuations (Kogan, Papanikolaou, Seru and Stoffman, 2017)–women display only a slight edge over their male counterparts. Hence, the considerable scarcity of women in innovation cannot solely be attributed to a labor market distortion that operates through positive selection.[3] Furthermore, the marginal difference in educational attainment between male and female R&D professionals doesn't support the possibility that a human capital distortion, which deters educational pursuit, could play a prominent role.

Taking the theory to the data, I find that lifting all barriers to female inventorship would increase U.S. income per person by 8.6% in the long run. Those economic gains are almost entirely achieved by raising exposure to innovation for aspiring women inventors. Indeed, to rationalize the disconcertingly skewed gender composition of U.S. inventors with the fact that women and men in innovation are just as productive and educated, I infer modest labor market and human capital distortions but large exposure distortions. Moreover, I show that the rise in income per person mostly comes from having *better* rather than *more* inventors.

From a social welfare perspective, eliminating distortions would be equivalent to permanently raising everyone's consumption by 2.7%. This figure comes shy of the long-run increase in income per capita due to slow transition dynamics. Of this improvement in welfare, 85% results from higher mean consumption while the remainder comes from lower consumption inequality. However, those gains are not evenly shared in the economy. Carrying out our consumption-equivalent welfare calculation separately for different demographic groups shows that future generations would experience a 3.6% permanent increase in consumption, as opposed to a more modest 0.4% increase for current cohorts. Zooming in on the current generation of inventors, women would see their consumption rise by 2% while men would instead suffer a 1.5% *decrease* therein.

The rest of the paper is outlined as follows. In the remainder of this section, we discuss the relevant literature. Section 2 presents the theoretical framework. Section 3 discusses the data on U.S. inventors and the model's calibration. Section 4 presents the results on the sources and macroeconomic implications of barriers to female innovation. Section 5 proceeds with various theoretical extensions and Section 6 concludes.

[3] This empirical fact further discards several extensions to the theory such as gender differences in inherited preferences for innovation or differential risk aversion between women and men when talent is partly unobserved, which would also operate through selection.

Related Literature

This paper relates and contributes to two growing strands of literature. The first is a collection of studies on the macroeconomic consequences of talent misallocation. A prominent example is Hsieh, Hurst, Jones and Klenow (2019) showing that convergence in the gender and racial composition of the U.S. labor market between 1960 and 2010 is responsible for as much as 40% of economic growth over that period. Hsieh and Moretti (2019) and Bryan and Morten (2019) study the allocation of talent across geographic locations and find that barriers to local migration can be a considerable drag on aggregate productivity. Lagakos and Waugh (2013) and Buera, Kaboski and Shin (2011) argue that selection on talent goes a long way in explaining the large productivity differences across countries. Morazzoni and Sy (2022) focus on entrepreneurs to document that financial frictions are particularly salient for women and that closing the gender gap in credit access could deliver sizable economic gains. Chiplunkar and Goldberg (2021) and Bento (2021) consider a larger set of barriers to female entrepreneurship and find similarly large productivity and welfare gains from eliminating gendered distortions.

Closer to this paper are studies that focus on the allocation of *inventive* talent. Celik (2022) argues that if inherited wealth is only weakly correlated with inventive ability, the overrepresentation of inventors from wealthy backgrounds is indicative of talent misallocation resulting from financial frictions. Akcigit, Pearce and Prato (2020) show that when aspiring inventors face financial barriers to human capital accumulation, education subsidies may be better suited than R&D tax credits in raising aggregate productivity. Lehr (2023) argues that firms' monopsony power over inventors can lead to a substantial misallocation of R&D. Arkolakis, Lee and Peters (2020) and Prato (2021) show that lifting immigration restrictions between Europe and the United States can reallocate inventors to where they are most productive and propel knowledge diffusion.

Closest to our paper is Einiö, Feng and Jaravel (2022) who first document that people from different social backgrounds and experiences produce innovations that are more tailored to their own needs. Unequal access to the innovation system can thus distort the direction of inventions, with potentially dire consequences for cost-of-living inequality and economic growth. To quantify the latter, the authors develop a two-sector endogenous growth model with heterogeneous consumer tastes and unequal access to innovation across different socio-demographic groups (including gender). Through the lens of this model, they find that barriers to female innovation are responsible for an 18.2% difference in the cost of living between women and men, and reduce the rate of economic growth by 1.4 percentage points. Our analyses are complementary as we extend their theoretical framework to further account for labor market discrimination and barriers to human capital accumulation for female inventors. Moreover, we propose a model of *semi*-endogenous growth with nontrivial transition dynamics due to the gradual accumulation of ideas and physical capital accumulation, and an overlapping generations structure. Hence, we find more modest effects of eliminating barriers to female innovation on productivity growth and welfare. Yet, our framework abstracts from the direction of innovation, which the authors show is an important margin in their counterfactual experiments. In that sense, our analyses provide complementary insights into the welfare consequences of unequal access to innovation by gender.

The second strand of literature to which this paper relates is the large body of empirical evidence on gendered barriers to innovation. Carrell, Page and West (2010) document that the findings of Bell et al. (2018) on the importance of having access to relevant role models persist further down the pipeline. Indeed, they show that female students are more likely to enroll in science and mathematics classes and go on to graduate with a STEM degree when they are assigned to female professors. However, Hunt, Garant, Herman and Munroe (2013) find that only 7% of the gender gap in patenting can be explained by women's lower probability of holding a STEM degree. Ross et al. (2022) provide evidence that women are 59% less likely to be credited with authorship on patents to which they contributed. When they do receive due credit, Jensen et al. (2018) find that their applications are more likely to be rejected, those rejections less likely to be appealed, and even for successful applications, women are granted a lower fraction of their claims, receive fewer citations and their patents are less likely to be maintained. This evidence is further supported by recent work from Hochberg et al. (2023) who use state-of-the-art tools from machine learning to estimate gender bias in patent citations. Kim and Moser (2021) document that, at the height of the baby boom, mothers who chose a career in innovation were extremely positively selected, patenting more than twice as much as women without children.

My contribution to this literature is threefold. First, I focus on a particularly large and salient source of talent misallocation: the underrepresentation of women among U.S. inventors. Women represent perhaps the largest pool of underutilized inventive talent, suggesting that there is ample scope to expand aggregate research effort. Second, I leverage detailed micro-level data on the universe of U.S. inventors to shed light on the sources of barriers to female innovation. Instead of honing in on a particular friction in isolation, I let the data speak through the lens of a unifying theory to compare those barriers in common units. Such "apples to apples" comparisons are indispensable to discriminate across competing theories and ultimately identify policy and research priorities. Finally, the rich yet tractable general equilibrium framework developed in this paper brings new insights to a largely empirical literature on female innovation. It plays the role of a laboratory through which one can study the aggregate and distributional implications of various policy interventions when price adjustments and transitional dynamics can play a first-order role.

2 Theory

In this section, I develop a theory of endogenous growth inspired by the seminal work of Romer (1990) and Jones (1995). Central to this theory is the role played by the allocation of inventive talent. The model combines several key ingredients: semi-endogenous growth through an expanding measure of varieties, an overlapping generations structure, physical and human capital accumulation, inventive talent heterogeneity, an irreversible occupation choice and, most importantly, gendered barriers to innovation.

The order of exposition proceeds as follows. In Section 2.1, I present the economic environment, devoid of distortions, laying bare the primitive driving forces of the model. Sections 2.2 and 2.3 then introduce gendered barriers, describe individual decision problems and characterize their aggregation. Finally, Section 2.4 defines the equilibrium allocation and Section 2.5 discusses the intuition underlying the model.

2.1 Economic Environment

Population and Preferences

The economy is populated by a measure N_t of working individuals indexed by i and their gender $g \in \{\text{women}, \text{men}\}$.[4] In the spirit of Yaari (1965) and Blanchard (1985), the working population features overlapping generations denoted by κ in which individuals retire at rate d. New individuals enter employment at rate b to form the most recent cohort such that the working population evolves as:

$$\dot{N}_t = nN_t \quad \text{where} \quad n = b - d > 0. \tag{1}$$

Individuals have logarithmic preferences over consumption c_{it} and isoelastic disutility from pre-career schooling time s_i such that lifetime utility is defined as:

$$U_i = \int_{\kappa}^{\infty} e^{-(\rho+d)(t-\kappa)} \ln(c_{it}) \mathrm{d}t - \beta s_i^{1+\nu} \tag{2}$$

where $\rho > 0$ is the pure rate of time preference, d reflects discounting from the stochastic retirement rate, $\beta > 0$ measures the strength of preferences over pre-career leisure time and $\nu > 0$ is the inverse of the Frisch elasticity of schooling time.

[4]The population is composed of as many women and men.

Technology

There are three sectors in the economy: the final, intermediate and research sectors. The final sector uses a variety of intermediate inputs indexed by j to produce a final good Y_t:

$$Y_t = \left(\int_0^{A_t} y_{jt}^{\frac{\sigma-1}{\sigma}} \mathrm{d}j \right)^{\frac{\sigma}{\sigma-1}} \tag{3}$$

where y_{jt} is the quantity of intermediate j used in production, A_t is the measure of existing varieties and $\sigma > 1$ is the elasticity of substitution between those varieties. Each intermediate input is produced by a single firm from the intermediate sector using physical capital k_{jt} and production labor ℓ_{jt}:

$$y_{jt} = k_{jt}^{\alpha} \ell_{jt}^{1-\alpha} \tag{4}$$

where $\alpha \in [0, 1]$ is the output elasticity of physical capital. However, before it can be brought to market, a variety must first be discovered. This role is played by the research sector, which combines research labor R_t with the existing stock of "ideas" A_t to develop varieties that are entirely new to society:

$$\dot{A}_t = A_t^{\phi} R_t. \tag{5}$$

As discussed in Jones (1995), ϕ measures the strength of knowledge spillovers. If ϕ is positive, the discoveries of yesterday make inventors more productive today. If instead ϕ is negative, it becomes harder and harder to find new ideas (Bloom, Jones, Van Reenen and Webb, 2020).

Occupations and Endowments

Individuals can be allocated to work in one of two occupations. A person can either supply production labor as a worker or research labor as an inventor. Each person is born with some *innate* inventive talent z_i, drawn from a Pareto distribution with cumulative distribution function G:

$$G(z) = 1 - z^{-\theta} \tag{6}$$

where the shape parameter $\theta > 1$ measures the degree of talent dispersion.

Upcoming cohorts must also allocate time s_i towards schooling to accumulate human capital h_i before starting their career:

$$h_i = s_i^{\eta} \tag{7}$$

where $\eta \in (0, 1)$ measures the return to schooling. Since everyone is endowed with a single unit of time in each period, a person either supplies h_i units of production labor as a worker or takes advantage of their talent to supply $z_i \times h_i$ units of research labor as an inventor. Both types of labor are supplied inelastically.

Resource Constraints

The final good can either be invested in physical capital or spent on consumption, which delivers the resource constraint:

$$\dot{K}_t + \delta K_t + C_t \leq Y_t \quad \text{where} \quad K_t \equiv \int_0^{A_t} k_{jt} \mathrm{d}j \quad \text{and} \quad C_t \equiv \int_0^{N_t} c_{it} \mathrm{d}i \tag{8}$$

and where $\delta > 0$ is the rate at which physical capital depreciates. The resource constraint for research labor is:

$$R_t \leq \int_0^{N_t} \mathbb{1}_{\{i \in R\}} z_i h_i \mathrm{d}i \tag{9}$$

and the resource constraint for production labor is:

$$L_t \leq \int_0^{N_t} \mathbb{1}_{\{i \in L\}} h_i \mathrm{d}i \quad \text{where} \quad L_t \equiv \int_0^{A_t} \ell_{jt} \mathrm{d}j. \tag{10}$$

The economic environment is summarized in Table 1.

2.2 Decision Problems

The Final Sector's Problem

The final sector is assumed to be perfectly competitive on the final good and intermediate input markets. Hence, it chooses how much of each variety to produce with as to maximize profits while taking the measure of varieties A_t and prices as given:

$$\max_{y_{jt}} \{P_t Y_t - \int_0^{A_t} p_{jt} y_{jt} \mathrm{d}j\}$$

Table 1: The Economic Environment

(1)	$\dot{N}_t = nN_t$	Population
(2)	$U_i = \int_\kappa^\infty e^{-(\rho+d)(t-\kappa)} \ln(c_{it})\mathrm{d}t - \beta s_i^{1+\nu}$	Lifetime utility
(3)	$Y_t = (\int_0^{A_t} y_{jt}^{\frac{\sigma-1}{\sigma}} \mathrm{d}j)^{\frac{\sigma}{\sigma-1}}$	Final good production
(4)	$y_{jt} = k_{jt}^\alpha \ell_{jt}^{1-\alpha}$	Variety production
(5)	$\dot{A}_t = A_t^\phi R_t$	Variety creation
(6)	$z_i \sim \text{Pareto}(\theta)$	Inventive talent
(7)	$h_i = s_i^\eta$	Human capital
(8)	$\dot{K}_t + \delta K_t + C_t \leq Y_t$	Final good resource constraint
(9)	$R_t \leq \int_0^{N_t} \mathbb{1}_{\{i \in R\}} z_i h_i \mathrm{d}i$	Research labor resource constraints
(10)	$L_t \leq \int_0^{N_t} \mathbb{1}_{\{i \in L\}} h_i \mathrm{d}i$	Production labor resource constraint

where p_{jt} is the price of intermediate input j and P_t is the price of the final good, which is normalized to unity. This delivers the following demand functions:

$$y_{jt} = Y_t / p_{jt}^\sigma.$$

The Intermediate Sector's Problem

To hold a claim on a variety's perpetual profits, an intermediate firm must first purchase its patent from the research sector through free-entry. Once that firm holds a patent, it engages in monopolistic competition on the market for intermediate inputs and perfect competition in the physical capital and production labor markets. That is, it chooses a price as well as physical capital and production labor to maximize profits π_{jt} while taking as given the demand for its variety, the rental rate of physical capital r_t and the wage paid to workers w_t^L:

$$\pi_{jt} = \max_{p_{jt}, k_{jt}, \ell_{jt}} \{p_{jt} y_{jt} - (r_t + \delta) k_{jt} - w_t^L \ell_{jt}\}.$$

Intermediate firms thus set their price to a constant markup μ above marginal cost:

$$p_{jt} = \mu \left(\frac{r_t + \delta}{\alpha}\right)^{\alpha} \left(\frac{w_t^L}{1-\alpha}\right)^{1-\alpha} \quad \text{where} \quad \mu \equiv \frac{\sigma}{\sigma - 1}$$

which implies that profits are perfectly symmetric across firms:

$$\pi_{jt} = \frac{Y_t}{\sigma A_t} \quad \forall j \in [0, A_t].$$

The Research Sector's Problem

Similarly, the research sector is assumed to be monopolistically competitive in the market for patents and perfectly competitive in the research labor market. Hence, it chooses a patent price q_t as well as research labor to maximize profits, while taking as given the wage w_t^R paid to inventors and the measure of varieties:

$$\max_{q_t, R_t} \{q_t A_t^{\phi} R_t - w_t^R R_t\}.$$

This delivers the following research labor market clearing condition:

$$w_t^R = q_t A_t^{\phi}.$$

With free-entry among patent buyers, the research sector sets the price of a patent to extract all possible rents from the commercialization of an invention, which corresponds to the present value of a variety's stream of future profits:

$$q_t = \int_t^{\infty} e^{-\int_t^{t'} r_\tau \mathrm{d}\tau} \pi_{t'} \mathrm{d}t'.$$

The Individual's Problem

Taking prices as given, the problem of an individual i is to select a career in which they will choose consumption and schooling to maximize lifetime utility:

$$U_i = \max_{c_{it}, s_i} \int_{\kappa}^{\infty} e^{-(\rho+d)(t-\kappa)} \ln(c_{it}) \mathrm{d}t - (1 + \tau_{g\kappa}^H) \beta s_i^{1+\nu}$$

subject to the flow budget constraint:

$$\dot{a}_{it} = \begin{cases} r_t a_{it} + (1 - \tau^L_{gt}) w^R_t z_i h_i - c_{it} & \text{if } i \in R \\ r_t a_{it} + w^L_t h_i - c_{it} & \text{if } i \in L. \end{cases}$$

Here, a_{it} is the financial wealth of individual i at time t with initial condition $a_{i\kappa} = 0$. As described in Section 1, the labor market and human capital distortions τ^L and τ^H are intended to reflect different forms of discrimination based on gender, which may vary across time. In particular, those distortions are normalized to zero within the production sector, which is roughly equivalent to defining the distortions in the innovation sector as *relative* to those that would be faced by production workers. This simple problem structure implies that optimal consumption is proportional to total wealth (financial and human) and satisfies the usual Euler equation:

$$c_{it} = \begin{cases} (\rho + d)[a_{it} + (1 - \tau^L_{gt}) \omega^R_t z_i h_i] & \text{if } i \in R \\ (\rho + d)(a_{it} + \omega^L_t h_i) & \text{if } i \in L \end{cases} \quad \text{and} \quad \frac{\dot{c}_{it}}{c_{it}} = r_t - \rho - d$$

where ω^o_t denotes the present value of the stream of future wages:

$$\omega^o_t \equiv \int_t^\infty e^{-\int_t^{t'} r_\tau \mathrm{d}\tau} w^o_{t'} \mathrm{d}t' \quad \forall o \in \{R, L\}.$$

Hence, at any point in time, a person expects current distortions to prevail for the rest of their career. In that sense, individuals have perfect foresight over all future prices *conditional* on current distortions remaining constant forever.

New cohorts receive no inheritance and choose how much time to allocate towards schooling before entering employment. Over time, individuals accumulate savings to eventually retire and consume their remaining financial wealth in retirement. Hence, a person's expected consumption in period t from the point of view of period κ is:

$$c_{it} = \begin{cases} (\rho + d)(1 - \tau^L_{g\kappa}) \omega^R_\kappa z_i h_i e^{\int_\kappa^t (r_\tau - \rho - d) \mathrm{d}\tau} & \text{if } i \in R \\ (\rho + d) \omega^L_\kappa h_i e^{\int_\kappa^t (r_\tau - \rho - d) \mathrm{d}\tau} & \text{if } i \in L. \end{cases}$$

Substituting this expression in the definition of lifetime utility and choosing schooling time to maximize it delivers:

$$s_i = \begin{cases} [\beta(1 + \nu)(\rho + d)(1 + \tau^H_{g\kappa})/\eta]^{\frac{-1}{1+\nu}} & \text{if } i \in R \\ [\beta(1 + \nu)(\rho + d)/\eta]^{\frac{-1}{1+\nu}} & \text{if } i \in L. \end{cases} \tag{11}$$

Schooling time is unsurprisingly decreasing in the human capital distortion, which raises disutility from education. Substituting this optimal choice back into the definition of lifetime utility as perceived from period κ, we have that:

$$U_i = \begin{cases} \ln[(1-\tau^L_{g\kappa})(1+\tau^H_{g\kappa})^{\frac{-\eta}{1+\nu}} \times \omega^R_\kappa \times z_i] + \text{const.} & \text{if } i \in R \\ \ln(\omega^L_\kappa) + \text{const.} & \text{if } i \in L. \end{cases}$$

Some individuals, irrespective of their talent, might never encounter the opportunity to become inventors. In this context, only those exposed to innovation can opt for it as a career path. Drawing from Bell, Chetty, Jaravel, Petkova and Van Reenen (2019), this exposure is modeled as a Bernoulli random variable e_i with mean parameter $1-\tau^E_{g\kappa} \in [0,1]$. The distortion $\tau^E_{g\kappa}$ reduces the proportion of individuals from cohort κ and gender g who can consider research as a career option. Conditional on receiving exposure to innovation, a person then decides whether to pursue research or work in production depending on which of those two occupations promises the highest lifetime utility. This occupation choice is made once and is irreversible thereafter. Thus, an individual will opt for research if and only if their inventive talent exceeds the threshold $\underline{z}_{g\kappa}$:

$$\underline{z}_{g\kappa} \equiv \underbrace{\frac{(1+\tau^H_{g\kappa})^{\frac{\eta}{1+\nu}}}{1-\tau^L_{g\kappa}}}_{\text{Distortions}} \times \underbrace{\frac{\omega^L_\kappa}{\omega^R_\kappa}}_{\text{Wages}}. \tag{12}$$

This selection threshold clearly illustrates how larger distortions effectively "raise the bar" differentially by gender, whereas a lower stream of *relative* future wages does so uniformly. That is, only the most talented people find it worthwhile to pursue research despite being paid below their marginal product, without adequate human capital, strictly because of their gender.

2.3 Aggregation

Having described individual decisions, let us now consider their aggregation. This section delves into how distortions impact both the quantity (extensive margin) and quality (intensive margin) of the overall research labor supply.

The Extensive Margin

The quantity of inventors by demographic group is obtained by aggregating individual occupation decisions. Accordingly, the fraction of individuals from gender g and cohort

κ who go on to invent is given by:

$$\mathrm{P}(z_i \geq \underline{z}_{g\kappa} \cap e_i = 1) = \underbrace{\frac{(1-\tau^E_{g\kappa})(1-\tau^L_{g\kappa})^\theta}{(1+\tau^H_{g\kappa})^{\frac{\theta\eta}{1+\nu}}}}_{\text{Distortions}} \times \underbrace{\left(\frac{\omega^R_\kappa}{\omega^L_\kappa}\right)^\theta}_{\text{Wages}} \tag{13}$$

This equation reflects the fact that a high-paying occupation will attract more people irrespective of gender, whereas distortions are solely responsible for gender differences in occupation choices. In contrast, the selection threshold $\underline{z}_{g\kappa}$ is entirely independent of exposure draws and instead strictly captures the trade-offs faced by someone deciding whether to pursue research rather than production. Notice that the fraction of inventors in a demographic group is inversely related to this threshold with elasticity θ. In fact, for large values of θ, the distribution of talent is tightly compressed around the selection threshold such that small changes in this cutoff induce large flows across occupations.

The Intensive Margin

Since career choices follow a cutoff rule on inventive talent, distortions may not only affect the *quantity* of inventors but also the *quality* of the resulting pool. Taking the product of talent and human capital and integrating over the resulting distribution delivers an expression for the average supply of research labor in cohort κ and gender g:

$$\mathbb{E}[z_i \times h_i | z_i \geq \underline{z}_{g\kappa} \cap e_i = 1] \propto \frac{\omega^L_\kappa}{(1-\tau^L_{g\kappa})\omega^R_\kappa}. \tag{14}$$

The average quality of inventors in that group is inversely proportional to the "keep rate" of the labor market tax. This reflects the selection mechanism. That is, if women inventors were only paid a fraction of their marginal product, only the very best would earn enough to prefer a career in research over one in production. Notice also that the human capital distortion does not appear in this expression. Indeed, more severe barriers to education give rise to positive selection, but they also directly hamper human capital formation, thus depressing research productivity. Altogether, the selection effect through talent and the direct effect through human capital exactly offset to leave the average productivity of inventors unchanged.

2.4 Equilibrium Allocation

Now that all decision problems, as well as their aggregation have been described, we can define the concept of an equilibrium allocation along a transition path.

Given time paths for distortions, an equilibrium consists of time paths for allocations and prices such that for all t:

1. $\{\{c_{it}, s_i\}_{i=0}^{N_t}\}_t$ and the occupation choice solve the individual's problem.
2. $\{\{y_{jt}\}_{j=0}^{A_t}\}_t$ solve the final sector's problem.
3. $\{\{p_{jt}, k_{jt}, \ell_{jt}\}_{j=0}^{A_t}\}_t$ solve the intermediate sector's problem.
4. $\{q_t, R_t\}_t$ solve the research sector's problem.
5. $\{A_t\}_t$ is given by equation (5).
6. $\{N_t\}_t$ is given by equation (1).
7. $\{K_t, Y_t, C_t\}_t$ satisfy the final good's resource constraint.
8. $\{\{p_{jt}\}_{j=0}^{A_t}\}_t$ clear the intermediate input markets.
9. $\{r_t\}_t$ clears the asset market: $\int_0^{N_t} a_{it}\mathrm{d}i = K_t + q_t A_t$.
10. $\{w_t^L, w_t^R\}_t$ clear the production and research labor markets.

2.5 Discussion

It is worth taking a step back at this point to discuss how distortions and parameters interact to drive the key mechanisms of our theory. To do so, let us start by reviewing the model's three sources of distortions. At the individual level, τ^L and τ^H distort the payoff to inventions differentially by gender, and since occupation choices follow a cutoff rule on talent, these distortions mostly influence the career choice of marginally talented individuals who are on the fence of the selection threshold. This is precisely where the sharp distinction with the exposure distortion resides. Echoing the arguments in Bell et al. (2018), a lack of exposure to inventive careers can deny opportunities to even the most talented aspiring inventors, which is what makes it particularly damaging.

At the aggregate level, all three sources of distortions either discourage or outright prevent women from pursuing innovation. However, as seen from equation (14), only the earnings tax distorts the quality of the resulting pool. Indeed, this tax effectively raises the talent threshold above which it becomes worthwhile to pursue innovation. This selection process raises research productivity by driving out the marginal inventors. The human capital distortion operates through that very same selection mechanism. However, it also directly depresses research productivity by discouraging *all* inventors from acquiring the skills and knowledge they need to innovate. Overall, these selection and direct effects exactly offset, thus leaving the quality of the inventive pool unchanged. Finally, since exposure to innovation is random and assumed to be uncorrelated with innate ability, it has no compositional effect through talent.

This discussion so far summarizes the qualitative implications of each distortion, but their quantitative relevance is mediated by the parameters of the model. Several parameters are particularly important: the cohort entry and exit rate b and d, the shape parameter θ of the talent distribution, the economy's overall degree of increasing returns to scale (here denoted by γ), and the strength of knowledge spillovers ϕ.

The Demographic Parameters.—b and d are most critical in quantifying the welfare implications of distortions. Since occupation choices are assumed to be irreversible, distortions follow individuals over their entire careers. Hence, one can think of b and d as effectively measuring the degree of churning in the economy, which directly influences the speed of transition dynamics. In other words, if old distorted cohorts "stick around" for too long, the economic gains from lifting distortions might only unfold in the distant future. Further, conditional on particular values of b or d, population growth ($n = b - d$) influences the rate at which those gains are discounted back to the present in that it tells us how many people will populate the future and enjoy the higher standards of living.

The Talent Parameters.—A fundamental insight of endogenous growth theory is that there exists an intertemporal trade-off in allocating resources between production and research. The former delivers higher consumption today, while the latter promises even more consumption tomorrow. In our framework, the key resource to allocate is people who differ in their research productivity. This talent heterogeneity matters in that it dictates the degree to which it is possible to "economize" on inventors. Put differently, if we only needed to allocate the brightest and most creative people toward research, there would be more workers left to produce existing varieties. This is precisely where the parameter θ comes into play, which measures the degree of dispersion in research productivity. For lower values of θ, the distribution of talent admits a fatter right tail such that eliminating barriers to female innovation entails the reallocation of just a few star researchers. If we only needed the Gerty Coris, Gertrude Elions and Jennifer Doudnas of this world to dedicate their extraordinary talent to research, there would be more people left to help turn their ideas into widely available goods and services.

Increasing Returns to Scale.—As elegantly described in Romer (1990), the nonrivalry of ideas gives rise to increasing returns to scale. More concretely, in any idea-based growth theory, income per person y_t, in the long run, is roughly proportional to the effective number of researchers R_t raised to some power γ:

$$y_t \propto R_t^{\gamma}$$

where $\gamma > 0$ measures the degree of increasing returns to scale. Intuitively, γ tells us the extent to which inventors matter as a propelling force for our living standards. To put it into perspective, if we were to double the number or quality of researchers, income per person would approximately rise by a factor of 2^{γ}. This simple example goes to illustrate how quantitatively critical is the degree of increasing returns to scale for our counterfactual exercises. In our theory, γ is a combination of parameters described by:

$$\gamma \equiv \frac{1}{(1-\alpha)(\sigma-1)(1-\phi)}.$$

Knowledge Spillovers.—Finally, even conditional on a particular value for the degree of increasing returns to scale, the knowledge spillover parameter ϕ has nontrivial implications for the dynamics of allocative efficiency. If inventors "stand on the shoulders of giants" ($\phi > 0$) but we somehow fail to take advantage of our current talent pool, not only do we miss out on new ideas that would improve all of our lives today, but we also prevent future generations of inventors from building on those ideas. This means that even if we could fix the allocation of talent in a snap of fingers, misallocation from the past could still cast a long shadow over our future economic well-being. In contrast, if it is harder and harder to find new ideas ($\phi < 0$) as emphasized in Bloom et al. (2020), that past misallocation could be a silver lining for future economic growth. Indeed, it would suggest that some of our best ideas are still out there waiting to be discovered. Atkeson, Burstein and Chatzikonstantinou (2019) and more recently Jones (2021) provide a more comprehensive analysis and discussion of the role of knowledge spillovers in transition dynamics.

3 Quantification

Let us now discuss how sensible empirical counterparts to the extensive and intensive margins of inventive effort, as well as educational attainment by gender can be used to quantitatively discipline the model's key driving forces and help us shed light on the barriers faced by women inventors. The data underlying the analysis is first described, after which I discuss the model's parameterization.

3.1 Data

To take our theory to the data, the first question we ought to ask ourselves is: who should be defined as an inventor? In practice, one is spoiled for choice from a plethora of conceivable definitions: scientists, engineers, entrepreneurs, GitHub developers, or even artificial intelligence algorithms. This choice is, however, far from obvious as it is notoriously hard to track inventions and their origins. With this caveat in mind, I follow a long tradition in the literature and focus on U.S. patent grantees, on whom a vast amount of information is publicly available.

Hence, the data used in our application comes from PatentsView, which contains information on all utility patents granted by the U.S. Patent & Trademark Office (USPTO) since 1976. PatentsView uses a series of algorithms to uniquely disambiguate inventors over time and, most importantly, predict their gender from their first name.[5] The sample I consider is restricted to inventors residing in the U.S. to whom gender is successfully attributed. With those restrictions, there remain about 1.7 million inventors to whom 3.7 million patents have been granted.

The Extensive Margin

Figure 1(a) shows that in 1976, 4% of inventors were women, and by 2020, that fraction had only inched up to 12%.[6] Although the representation of women in U.S. inventorship did increase by a factor of three over the last 40 years, it remains disconcertingly low. As a reference point, Figure 1(b) plots the same series but for lawyers, doctors, and engineers in the U.S. Censuses and American Community Surveys (ACS). Women respectively composed 6% and 11% of lawyers and doctors in 1970, with those figures both climbing to about 40% in 2020. In contrast, there has been much less convergence in the gender composition of engineers, where women were and are still vastly underrepresented at 2% of the profession in 1970 and only 15% in 2020.

Figure 2 plots the previous series, but by cohort instead of year.[7] Here, an inventor's cohort is defined as the year of first appearance in the data, and since the PatentsView data starts in 1976, cohort inference is restricted to inventors who first appear after 1986. This figure shows that convergence is slowly but clearly underway as new cohorts composed of more women gradually replace previous ones.

However, as evident in Figure 3, there is a nontrivial degree of heterogeneity in the

[5]For more information on the gender attribution algorithm used by PatentsView, see https://patentsview.org/gender-attribution.

[6]In this plot and hereafter, years correspond to the date at which a patent is granted.

[7]Cohorts are grouped in 5-year periods indexed by their ending year.

Figure 1: Representation of Women in Different Occupations

(a) Inventors

(b) Lawyers, Doctors and Engineers

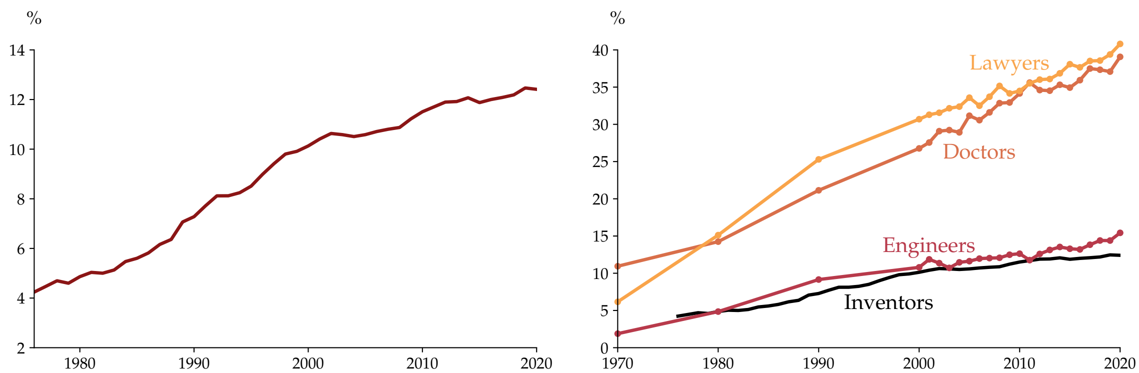

Note: The share of women among U.S. inventors has been slowly but steadily increasing from 4% in 1976 to 12% in 2020. This path resonates with the experience of women engineers in those same years but is in sharp contrast with the much faster convergence that occurred in the legal and medical professional spheres. Author's calculation from the U.S. Censuses and ACS.

Figure 2: Female Share of Inventors by Cohort

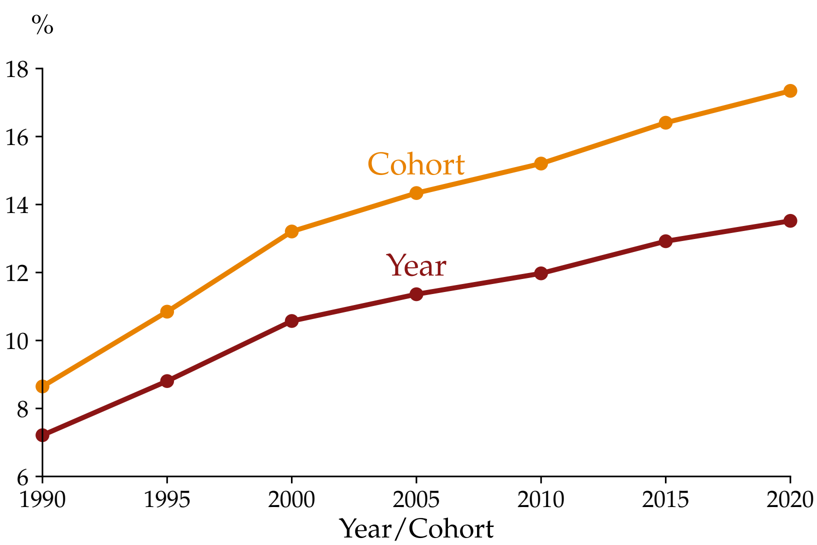

Note: The fraction of women in the 1990 cohort of U.S. inventors was 7.5% and now stands at 17.6% among the 2020 cohort. Cohorts are grouped in 5-year periods indexed by their ending year.

gender composition of different technological fields.[8] In that figure, technological fields

[8] Those broad fields are defined according to the single digit classes of the Cooperative Patent Classification (CPC) system, with the exception that the "textiles and paper" and "fixed constructions" classes are

are divided in two groups, as shown in panels 3(a) and 3(b): fields where women have been relatively more and less represented over our sample period. To highlight the two most contrasting examples, between the 1990 and 2020 cohorts, the female share of inventors went from 13% to 28% in the field of chemistry and metallurgy, and from 5% to 10% in the field of fixed constructions and mechanical engineering. Section 5 quantitatively explores the role of this heterogeneity across technological fields.

Figure 3: Female Share of Inventors by Technological Field

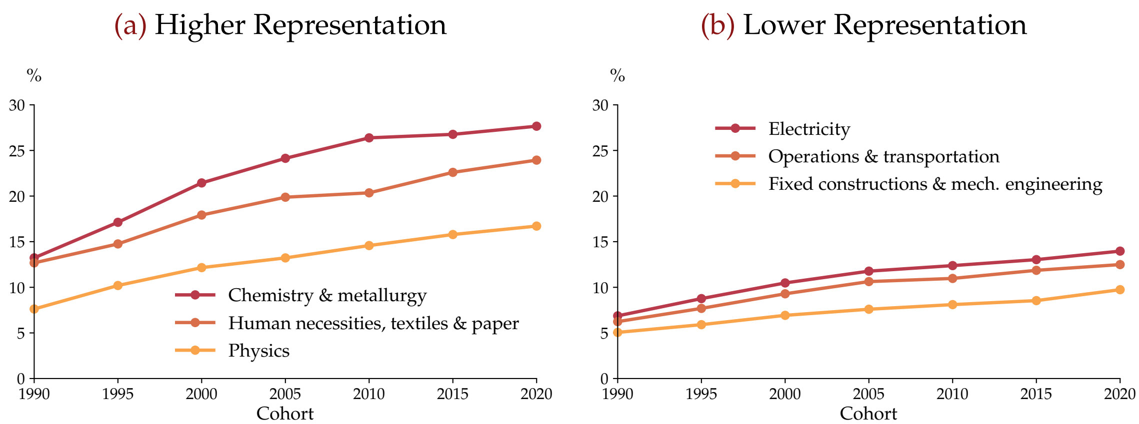

Note: Between the 1990 and 2020 cohorts, the female share of inventors in the field of chemistry and metallurgy has grown from 13% to 28%. In contrast, it went from 5% to 10% in the field of fixed constructions and mechanical engineering.

The Intensive Margin

As mentioned above, a notable feature of the PatentsView data is its state-of-the-art disambiguation procedure. It allows us to uniquely identify inventors across multiple patent issues to measure the number of patents granted to each of them every year. However, simple patent counts might only deliver a partial depiction of individual inventive productivity, as it has long been recognized that not all patents are invented equal (Griliches, 1990). Then, which proxy of a patent's quality should we consider? A familiar candidate is the number of citations received by a patent, but one might raise doubts on whether this is the right choice in the context of a study focused on gender. After all, citations are deliberately chosen by applicants and examiners who are not insusceptible to their own gender biases.

respectively grouped with the "human necessities" and "mechanical engineering" ones.[9] In particular, an inventor's field is here defined as the modal CPC class across all patents they have been granted over their career.

A plausible instance of such biases is documented by Jensen et al. (2018) looking at U.S. patents by lone inventors. They show that among inventors with relatively common forenames (where gender is easily inferred), women are cited 30% *less* frequently than men. Conversely, among inventors with rare first names, this pattern is completely reversed, with women being cited 20% *more* frequently than men.[10] A caveat of this exercise is that it could, in principle, reflect differential gender selection into inventorship for domestic and foreign applicants. Yet, citations are not the only margin on which women might be denied credit for their scientific contributions. Indeed, Ross et al. (2022) show that women are 59 percent less likely to be attributed authorship on patents to which they contributed. This suggests that the PatentsView data might be missing a nontrivial fraction of the female inventor population.

To better gauge the value of each patent, I instead weigh them by their economic significance as reflected in stock market reactions to patent grant announcements.[11] This valuation approach, introduced by Kogan et al. (2017), offers an economically sound perspective on the valuation of innovations, bridging scientific significance with market relevance. Critically, by grounding comparisons in stock market valuations rather than subjective interpretations, it offers a safeguard against potential biases and thus yields a more impartial evaluation of a patent's worth. To further refine this valuation, I adjust the measure for systematic valuation differences across fields using 3-digit CPC class fixed effects and account for the individual contribution of an inventor by controlling for co-inventorship team size.

Next, we must aggregate value-weighted patents at the level of inventors. To do so, we first take the sum of all value-weighted patents for each inventor in each year, conditioning on differences in patenting rates across 3-digit CPC classes. This delivers a measure of annual inventive output for each inventor in our sample and for all years in which they appear in the data. Then, to condense those multi-year observations to a single number, we first purge them from experience fixed effects, where experience $=$ year $-$ cohort. This is consistent with the model being silent on the life-cycle of inventive productivity. Then, we take the average across years for each inventor to obtain the key empirical counterpart to individual inventive productivity in our theory.

Figure 4 plots the distribution of individual inventive output by gender from two angles. The histogram depicted in Panel 4(a) illustrates the striking similarity in the distribution of inventive productivity between genders. Meanwhile, Panel 4(b) offers a complementary perspective by charting the female share of top inventors by cohort for

[10]These estimates not only condition on technological fields but also on forename frequency to control for any association between the rarity of an inventor's name and citations received.

[11]We find similar results using either truncation-adjusted forward citations or the notion of a patent's importance as proposed by Kelly, Papanikolaou, Seru and Taddy (2021).

different productivity quantiles. Although women are slightly overrepresented among the upper echelons of inventors, the discrepancies in representation across productivity tiers are remarkably small.

Figure 4: Distribution of Individual Inventive Productivity by Gender

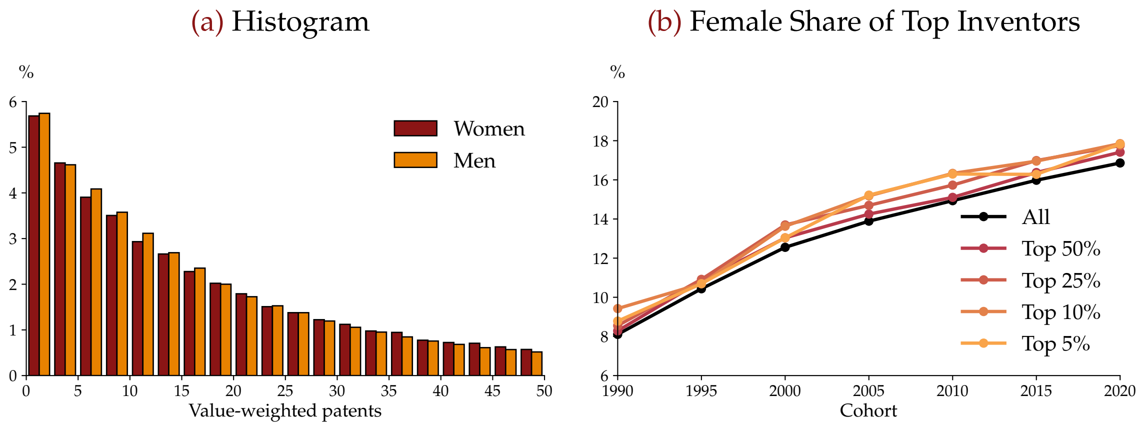

Note: As evident in these panels, women are only slightly overrepresented among top inventors.

As depicted in Figure 5, female inventors only exhibit a marginal productivity edge over their male counterparts. To be more specific, this figure plots the gender gap in inventive output–with a positive gap revealing a productivity *advantage* by women–controlling for technological fields.[12]

Educational Attainment

Earlier discussions highlighted the implications our model carries concerning the relative educational attainment of female and male innovators. However, a shortcoming of the PatentsView data is its lack of details on inventors' educational backgrounds. To the best of my knowledge, the most notable attempt at bridging this gap using publicly available data is from Hunt et al. (2013). They use the 2003 wave of the National Survey of College Graduates (NSCG), which recorded whether respondents were granted a patent in the last five years, and find that only 7% of the gender gap in patenting can be accounted for by women's lower propensity to hold a STEM degree.[13] Instead, they show that 78% and 15% of the gap are due to gender differences in patenting among those with and without a STEM degree, respectively.

[12]Notably, when analyzing these outcomes across individual technological domains, the variations are remarkably subtle.

[13]The survey only asked this question in the 1995 and 2003 waves.

Figure 5: Inventive Productivity Gender Gap

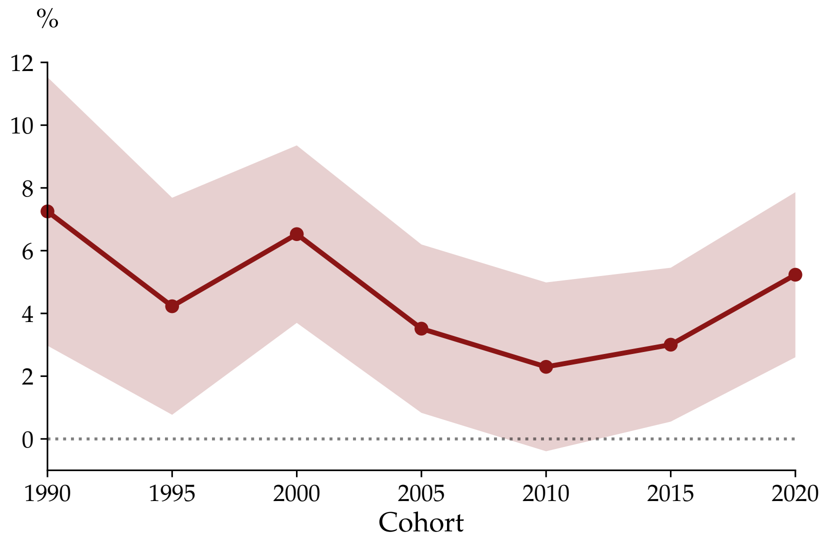

Note: Female inventors exhibit a modestly higher productivity relative to their male colleagues. The shaded area corresponds to the 95% confidence interval around the estimates.

To approximate the educational attainment gender gap among innovators over a more extensive time range, I use data from the NSCG between 1993 and 2019. More specifically, I restrict the sample to employed individuals between 30 and 34 years old whose main work activity is R&D.[14] Figure 6 plots the fraction of R&D workers in three educational attainment tiers over time for women and men. The standout observation here is how strikingly similar the two distributions are. Between 1993 and 2019, women and men in R&D consistently seemed to achieve the same educational attainment.

3.2 Calibration

Our model features eighteen parameters to be determined, of which three are set to standard values in the literature. The pure rate of time preference ρ is set to 0.02, the Cobb-Douglas parameter for physical capital α is set to $1/3$, and a 5% depreciation rate δ is assumed. Turning to the demographic parameters, the retirement rate d is set to $1/30$ to match an expected work-life of 30 years (reflecting the prime-age interval of 25 to 54 years old), and the cohort arrival rate b is correspondingly chosen to achieve a

[14] As mentioned above, in the 1995 and 2003 waves of the survey, respondents were asked about their patenting record and 71% and 47% of patentees reported R&D as their main work activity, respectively. Moreover, in those same two waves, R&D workers were 4.6 and 5.3 times more likely to patent than the average respondent.

Figure 6: Educational Attainment of R&D Workers by Gender

(a) Women (b) Men

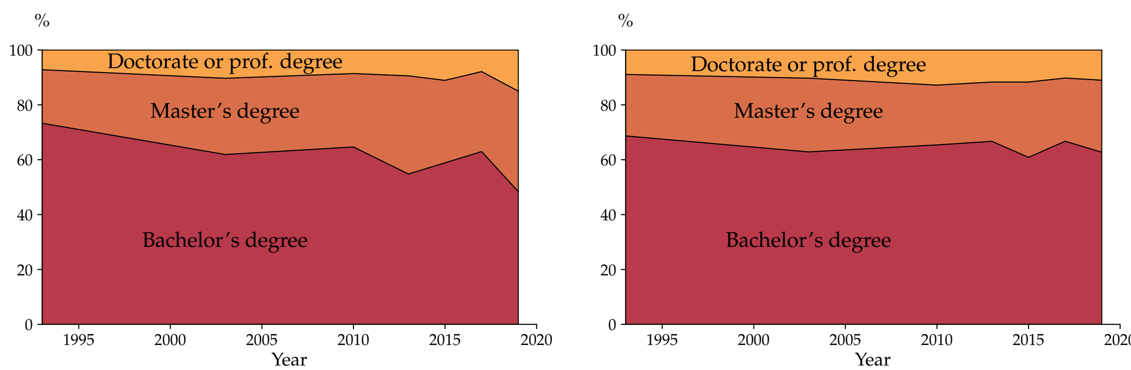

Note: This figure plots the fraction of employed R&D workers between 30 and 34 years old in each educational attainment category by gender and year. Author's calculation from the NSCG.

population growth of 0.5% per year, as projected by the U.S. Census Bureau until 2060.

To calibrate the human capital parameters, I assume that the pre-career period spans 25 years, which can either be spent on leisure or schooling. The schooling disutility parameter β is therefore set to 46.8 to match the 15.9 years of education that could be expected in the U.S. between 1990 and 2019 from the United Nations Development Program. The parameter η is chosen to approximate a Mincerian return to schooling of 10% around expected years of education, which is somewhere in the middle of the range of values reported in Card (1999).

We saw in Section 2.3 that the distribution of individual research labor supply follows a power law with tail exponent θ. Under perfectly competitive labor markets, individual research earnings also follow a power law with the same tail exponent. Hence, one can hope to recover θ from the empirical tail exponent of either of those two distributions. On one end, Bell et al. (2019) link patent records to tax records from 1999 to 2012 and estimate a tail exponent of 1.26 for the inventor earnings distribution. On the other, estimating the productivity distribution's tail exponent from value-weighted patents following the methodology of Clauset, Shalizi and Newman (2009) delivers a larger figure of 3.25.[15] In light of this uncertainty, I assume an intermediate value of $\theta = 2$, but Section 4.3 shows robustness to a relatively wide range of alternative values.

Finally, the knowledge externality parameter ϕ is set to -2.1, corresponding to the estimate of Bloom et al. (2020) for the aggregate U.S. economy, and the elasticity of substitution across varieties σ is set to 2.45 to target a degree of increasing returns to

[15]This is implemented with the powerlaw Python package of Alstott, Bullmore and Plenz (2014).

scale of one third, as suggested by Jones (2021). Notice that the chosen value for ϕ embraces the view that "ideas are getting harder to find". The calibration of the model is summarized in Table 2.

Table 2: Calibration

Parameter	Symbol	Value	Source
Discount rate	ρ	0.02	Standard calibration
Cobb-Douglas	α	1/3	Standard calibration
Depreciation rate	δ	0.05	Standard calibration
Retirement rate	d	1/30	Average work-life of 30 years
Entry rate	b	$d + 0.05$	0.5% U.S. population growth
Schooling disutility	β	46.8	15.9 expected years of schooling
Return to schooling	η	1.59	10% Mincerian return
Pareto shape	θ	2	Assumption
Knowledge spillover	ϕ	-2.1	Bloom et al. (2020)
Variety substitution	σ	2.45	Degree of IRS $= 1/3$

4 Results

We now have all the necessary ingredients to answer the two central questions of the paper: (1) What are the most salient barriers faced by women in innovation? and (2) How costly is the resulting misallocation of inventive talent for aggregate productivity and welfare? In Section 4.1, I first explain how moments in the data can be projected through the lens of our model to infer its three sources of distortions, and in Section 4.2, I quantify the macroeconomic implications of eliminating all barriers to female innovation.

4.1 Inferring Distortions

To infer the model's three sources of distortions, I leverage *cohort-level* moments on (1) the intensive and extensive margins of inventive effort by gender and (2) the educational attainment gender gap among U.S. R&D workers.

Labor Market Distortion

Looking back at Section 2.3, we saw from equation (14) that the average supply of research labor from cohort κ and gender g is inversely proportional to the "keep rate" of the corresponding labor market tax:

$$\mathbb{E}[z_i \times h_i | z_i \geq \underline{z}_{g\kappa}] \propto \frac{1}{1 - \tau^L_{g\kappa}}. \tag{15}$$

Equation (15) then implies that the labor market distortion can be directly inferred from the gender gap in average inventive output by cohort, as plotted in Figure 5. Panel 7(a) plots the inferred labor market tax, where two observations emerge. First, its magnitude is relatively subdued across all cohorts, which suggests that the large representation gender gap documented in Figure 1(a) cannot be rationalized by this distortion alone. Second, the labor market tax has been declining over time, although modestly, which is consistent with evidence from other occupations (Hsieh et al., 2019).

Human Capital Distortion

Equation (11) of Section 2.2 implies that the fraction of time spent on schooling by someone from cohort κ and gender g is also inversely proportional to the corresponding human capital distortion:

$$s_i \propto \frac{1}{1 + \tau^H_{g\kappa}}. \tag{16}$$

Therefore, if we were to observe inventors' schooling trajectory, we could readily infer human capital distortions from the educational attainment gender gap. Unfortunately, as discussed earlier, such data is not publicly available. Hence, to approximate the educational attainment gender gap among innovators, I instead use data from the NSCG between 1993 and 2019. To translate educational attainment categories to completed years of education, I assume that a bachelor's, master's, and professional or doctorate degree is equivalent to 16, 18, and 22 years of education, respectively. The derived human capital distortion can be seen in Panel 7(a). Similar to the labor market distortion, the magnitude of the human capital distortion is relatively minor, with female and male R&D professionals attaining comparable educational levels throughout the sample.

Figure 7: Inferred Distortions

(a) Labor Market and Human Capital (b) Exposure

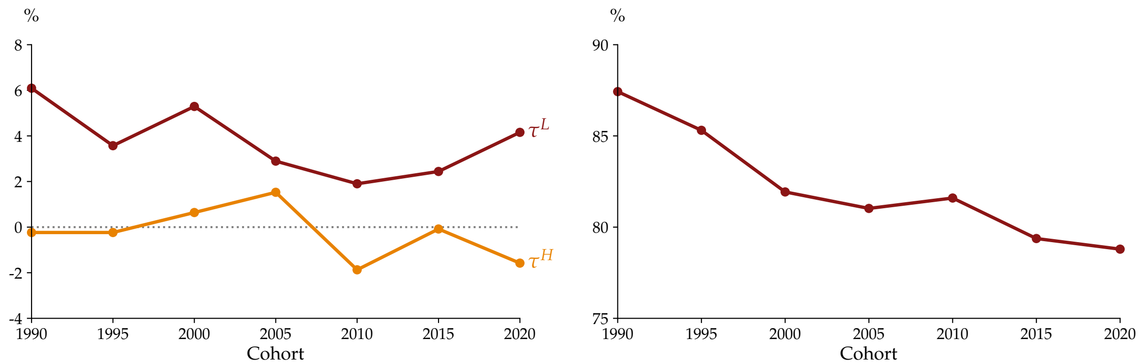

Note: Since women and men inventors are similarly productive and educated, I infer modest labor market and human capital distortions. In contrast, the exposure distortion is considerable, but unwaveringly declining over time.

Exposure Distortion

Given other distortions and parameter values, the exposure barrier is inferred from the gender ratio in the proportion of inventors. As we saw in Section 2.3, the fraction of inventors among gender g and cohort κ is proportional to:

$$P(z_i \geq \underline{z}_{g\kappa} \cap e_i = 1) \propto \frac{(1-\tau^E_{g\kappa})(1-\tau^L_{g\kappa})^\theta}{(1+\tau^H_{g\kappa})^{\frac{\theta\eta}{1+\nu}}}$$

The empirical counterpart to that fraction is calculated as the total number of inventors from gender g and cohort κ in the PatentsView data, divided by the labor force of gender g aged 25 to 35 in year κ of the Current Population Survey (CPS). This denominator accounts for gender differences in labor force participation during our sample period, on which our theory is silent.

Figure 7(b) plots the inferred exposure distortion faced by aspiring women inventors of each cohort. This distortion is strikingly large and stands in sharp contrast with the previous two. But this magnitude is perhaps not so surprising. Since female and male inventors are just as productive and educated, only a large barrier that does not operate through selection or human capital can rationalize the vast underrepresentation of women in research.

4.2 Eliminating Distortions

How "costly" are the barriers to female innovation and the resulting misallocation of inventive talent for U.S. aggregate productivity and welfare? To answer this question, a natural counterfactual exercise is to lift all distortions faced by women inventors starting in 2020 and compare the economy's transition path to its initial balanced growth path allocation.[16] Note that this starting point is consistent with the distortions inferred for the most recent cohort in Section 4.1.

Aggregate Productivity

Panel 8(a) shows that a simultaneous withdrawal of all distortions would raise income per person by 8.6% in the long run. Although relatively large, this productivity gain is very slow to materialize, with a half-life of a little under a century. This inertia is consistent with the findings of Atkeson et al. (2019) who study the transitional dynamics of a large class of semi-endogenous growth models with a gradual accumulation of physical capital and ideas. On top of these two forces, our theory further admits an overlapping generations structure with irreversible occupation choices, which directly slows down the reallocation of labor.

This labor reallocation is depicted in Panel 8(b), which plots the path of aggregate research labor by gender in percentage deviation from their starting point. The supply of research labor by women more than triples, while that of men shrinks by almost a quarter. In aggregate, research labor permanently increases by 28% within the first 50 years of the transition. As discussed in Section 2.5, raising this factor of 1.28 to the power of $\gamma = 1/3$ (the degree of increasing returns to scale) approximately recovers our long-run gain in income per person of 8.6%.

However, is this gain achieved by having *more* or *better* inventors? Figure 9 reveals the answer by tracing the path of the extensive and intensive margins of research labor. More specifically, Panel 9(a) plots the fraction of inventors in each gender group. About 0.47% of women are inventors in 2020, with that fraction ultimately rising to 1.22% after 250 years.[17] Marginally talented men are instead being gradually pushed out of

[16] The model's transition path is solved using the Relaxation Algorithm developed by Trimborn, Koch and Steger (2008). For additional information on its implementation, see https://sites.google.com/view/relaxmacro.

[17] The average share of researchers in the U.S. stood at almost 1% in 2020, which is not too far from what is suggested by the model even though this moment is not explicitly targeted in the calibration

Figure 8: Income Per Person and Aggregate Research Labor

(a) Income Per Person (b) Aggregate Research Labor

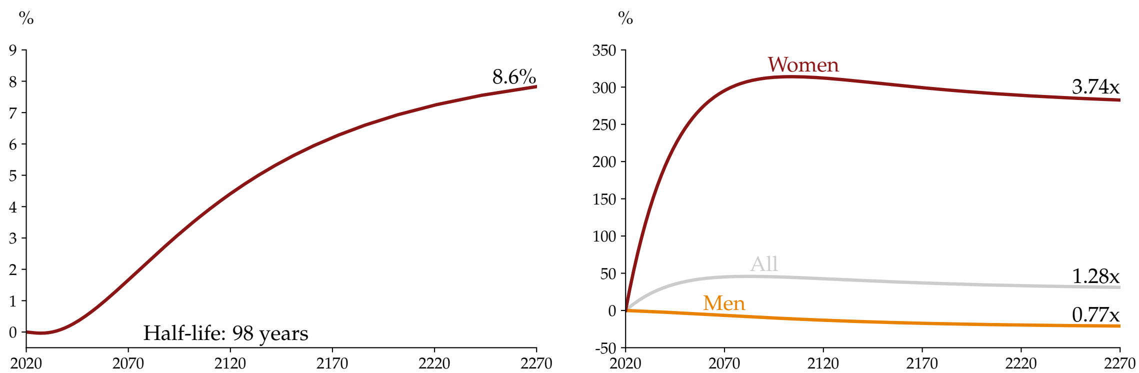

Note: Removing all barriers to female innovation could increase income per capita growth by about 4 basis points over the next century.

inventorship as older cohorts retire, culminating in a balanced gender representation across professions. Interestingly, the aggregate share of inventors barely rises in the long run, suggesting that most of the productivity gains from lifting distortions are achieved on the intensive margin.

Figure 9: The Extensive and Intensive Margins of Aggregate Research Labor

(a) Extensive Margin (b) Intensive Margin

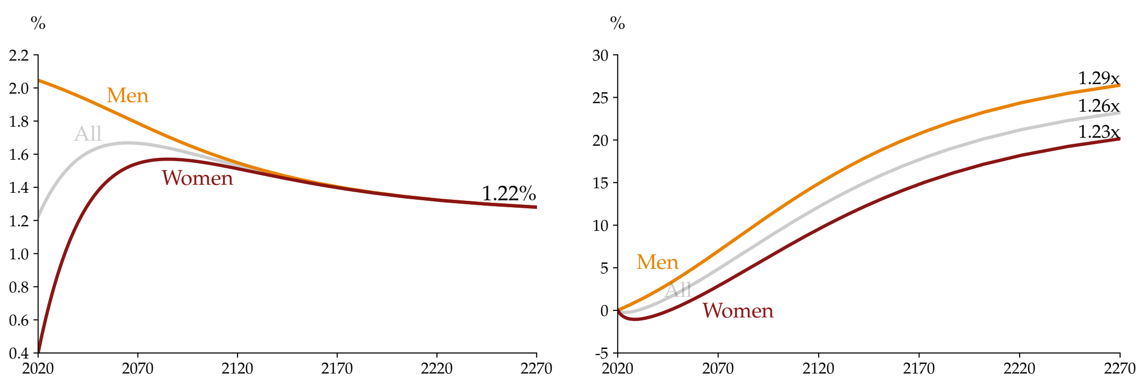

Note: The number of women inventors almost triples, while it shrinks by almost 40% for men. However, both gender groups become more productive on average as they face more intense competition from a larger pool of talented women.

(see https://data.oecd.org/rd/researchers.htm). According to the 2015 Frascati Manual, researchers are defined as "professionals engaged in the conception or creation of new knowledge, products, processes, methods and systems, as well as in the management of the projects concerned".

To drive this point home, Panel 9(b) plots the transition path of average inventive productivity by gender, in percentage deviation from its starting point. This figure leaves no doubt that, in aggregate, the large increase in research labor is unfolding through the intensive rather than the extensive margin. However, a closer look at the dynamics *within* gender groups reveals more nuance. For women, only 16% of the near 4-fold rise in research labor is coming from the intensive margin. In comparison, the decline in male research labor is entirely driven by their shrinking numbers, despite the counteracting rise in their average productivity from tougher competition.

Welfare

To assess the welfare implications of eliminating all distortions, define the following utilitarian social welfare function as in Calvo and Obstfeld (1988):

$$W_t(\lambda) = \int_t^{\infty} e^{-(\rho-n)(\tau-t)} \int_{-\infty}^{\tau} be^{-b(\tau-\kappa)} \mathbb{E}[\ln(\lambda \times c_{i\tau})] \mathrm{d}\kappa \mathrm{d}\tau$$

where the term involving pre-career schooling is ignored as it delivers the same disutility to everyone, regardless of distortions. In this expression, the fraction of people from cohort κ at time τ is given by $be^{-b(\tau-\kappa)}$, the expectation is taken over individuals within cohorts, and $\lambda > 0$ permanently multiplies the consumption of every person. Note that by changing the order of integration, we can additively separate the social welfare of surviving and future cohorts:

$$\begin{aligned} W_t(\lambda) = & \int_{-\infty}^{t} be^{-b(t-\kappa)} \int_t^{\infty} e^{-(\rho+d)(\tau-t)} \mathbb{E}[\ln(\lambda \times c_{i\tau})] \mathrm{d}\tau \mathrm{d}\kappa && \text{Surviving cohorts} \\ & + \int_t^{\infty} be^{-(\rho-n)(\kappa-t)} \int_{\kappa}^{\infty} e^{-(\rho+d)(\tau-\kappa)} \mathbb{E}[\ln(\lambda \times c_{i\tau})] \mathrm{d}\tau \mathrm{d}\kappa && \text{Future cohorts.} \end{aligned}$$

With these definitions, we can ask: by what factor λ must we permanently adjust the consumption of everyone in a distorted economy to leave them as well off as if they spent the rest of their lives in an undistorted economy? The answer satisfies:

$$W_t(\lambda) = W_t^*(1)$$

where W_t and W_t^* denote social welfare in the distorted and undistorted economies, respectively. More precisely, the distorted economy is and remains on its balanced growth path, characterized by the distortions inferred for the 2020 cohort. Instead, the undistorted economy starts from that same initial point but is launched on a transition path thereafter.

This exercise reveals that removing barriers to female innovation would be equivalent to permanently raising everyone's consumption by 2.7%. That figure is notably lower than the long-run income per person gain of 8.6% for two reasons. First, the transition is very slow, meaning that the higher standards of living will mostly materialize in a remote future. Second, because U.S. population growth is projected to be quite slow as well, those future gains must be discounted back to the present at a relatively high rate.

Of this 2.7% consumption-equivalent welfare variation, 85% comes from higher mean consumption, while the rest comes from lower consumption inequality. This does not imply, however, that those gains are evenly shared in the economy. Indeed, carrying out this welfare calculation separately for different demographic groups shows that removing all distortions is equivalent to a 3.6% increase in consumption for future cohorts, as opposed to a more modest 0.4% increase for surviving cohorts.

Notably, the current generation of female inventors stands to gain significantly, with a potential 2% increase in their consumption-equivalent welfare. On the flip side, not everyone benefits: surviving cohorts of male inventors would face a 1.5% decline in their consumption. Those distributional consequences should remind us that when the costs of an intervention are concentrated and borne today, while its benefits are diffuse and materialize tomorrow, we ought to think carefully about its implementability.

4.3 Robustness to Parameter Values

In this section, our main findings are revisited with alternative parameter values. In particular, I discuss their robustness to the degree of increasing returns to scale γ, the shape parameter θ of the inventive talent distribution, the demographic parameters b and d, and the knowledge spillover parameter ϕ. Note here that for consistency, I re-calculate the exposure distortion when varying parameter values.

As discussed in Jones (2021), the overall degree of increasing returns to scale in the economy is both notoriously challenging to measure and fundamental in providing practical answers to some of our most pressing macroeconomic questions. Despite the sparsity of empirical work on the matter, there have been some valuable quantification attempts. Jones (2002) estimates values ranging from about 0.05 to 0.33 through a time-series econometric analysis. Peters (2021) instead leverages the pseudo-random resettlement of 8 million ethnic Germans into West Germany after the Second World War to estimate a value of nearly 0.6. Given this wide range of estimates, Panel 10(a) plots the long-run percentage gain in income per capita after eliminating all distortions, which ranges from about 1% to almost 16% for γ going from 0.05 to 0.6.

The research productivity shape parameter θ is similarly relevant and challenging

Figure 10: Long-Run Gain in Income per Person: The Role of γ and θ

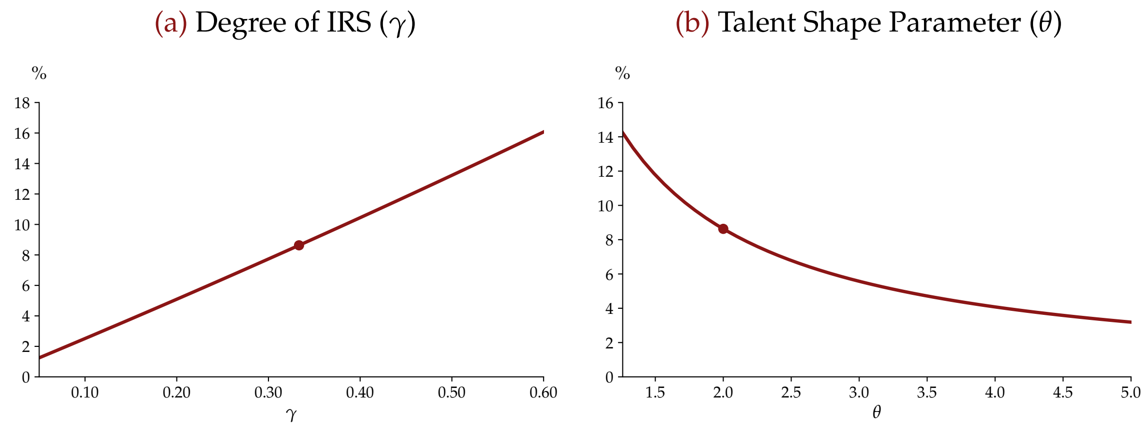

Note: The dots correspond to the long-run gain in income per person of 8.6% from our baseline calibration with $\gamma = 1/3$ and $\theta = 2$.

to determine accurately. Indeed, it is quantitatively relevant in that it determines the extent to which one can "economize" on the number of inventors by taking advantage of the extraordinary talent of a few star researchers. However, several obstacles cloud its precise estimation: (1) directly gauging individual research productivity presents a complex endeavor, and (2) inventor wages might not be dictated by perfect competition (Lehr, 2023). Nevertheless, Panel 10(b) reveals that the long-run productivity gain from lifting all distortions ranges from a little over 14% to about 3% when θ goes from 1.26 to 5. The lower bound of this range is borrowed from Bell et al. (2019), who estimate a tail exponent of 1.26 for the inventor earnings distribution between 1999 and 2012. The upper bound is instead chosen to be slightly larger than the tail exponent of the research productivity distribution (as measured from value-weighted patents), which I estimate to 3.25 following the methodology of Clauset et al. (2009).

The last set of parameters to which I assess robustness are those that speed up or slow down the transition toward a new balanced growth path. The demographic parameters b and d play precisely this role by dictating the cohort turnover rate. The knowledge spillover parameter ϕ does so by disciplining the degree of autocorrelation in the stock of ideas. In Table 3, I set d to values that correspond to an expected working life of 20 and 40 years and calculate the transition's half-life and the welfare gain from eliminating distortions.[18] Interestingly, the demographic parameters have virtually no effect on long-run income per capita. However, letting the expected working life vary by 10 years

[18] In particular, when varying d, I vary b by the same amount to keep the population growth rate constant, and I re-calibrate β to keep expected years of schooling unchanged. This isolates the counterfactual calculation from variations in the social discount rate and average human capital.

shifts the transition's half-life by roughly the same number of years. In present value terms, this translates to counterfactual welfare gains of 2.9% or 2.5% instead of our baseline figure of 2.7%.

Table 3: Income per Person, Half-Lives and Welfare: The Role of d and ϕ

Parameter	Value	Income Per Person Gain (%)	Half-Life (years)	Welfare Gain (%)
Baseline		8.63	97.9	2.67
d	1/20	8.63	87.1	2.91
	1/40	8.63	107.1	2.49
ϕ	0.8	8.63	112.6	2.26
	-6.2	8.63	96.9	2.70

Note: This table shows the long run percentage gain in income per person, its half-life and the percentage welfare gain from removing all distortions in 2020 for different values of d and ϕ. When varying d and ϕ, I respectively adjust b and σ as to keep population growth n, expected schooling and the degree of increasing returns to scale γ constant. The values of $1/20$ and $1/40$ for d correspond to an expected working life of 20 and 40 years, respectively. The values of 0.8 and -6.2 for ϕ span the range of values considered in Bloom et al. (2020).

For the knowledge spillover parameter ϕ, I select alternative values of 0.8 and -6.2, which span the values considered in Bloom et al. (2020). Importantly, here, when varying ϕ, I adjust σ to keep the degree of increasing returns to scale constant to $\gamma = 1/3$. That is precisely why variations in ϕ lead to no changes in long-run living standards in Table 3. Although lower values for the knowledge spillover parameter seem to barely change our results, a higher value of $\phi = 0.8$ raises the transition's half-life by about 15 years, which shrinks welfare gains by 41 basis points.

5 Theoretical Extensions

In which directions could we extend our model, and how would these refinements influence our counterfactual calculations? This section is an attempt to answer this question with seemingly important ingredients that were left out of the model. The derivations corresponding to these extensions are all presented in Appendix A.7.

Role Models and Affirmative Action

Previously, we noted that the exposure distortion was intended to encapsulate barriers such as the relative scarcity of relevant role models for young girls compared to young boys. What if we were to delineate this more explicitly within the model? By doing so, part of what is currently captured under the exposure distortion would be redefined as a technological friction, emphasizing the importance of role models in molding career aspirations. Such friction would have to be modeled as a technological constraint on the environment. Hence, achieving a more efficient talent allocation could be a longer journey, as women would only gradually join the ranks of inventors, in turn inspiring the next generation to do so.

To incorporate the influence of role models in the model, an individual's exposure to innovation is still represented as a Bernoulli random variable. However, its mean parameter now also depends on the proportion of individuals from each gender who opted for research in prior generations:

$$e_i \sim \text{Bernoulli}((1-\tau^E_{g\kappa}) \times I^{\epsilon_g}_{g\kappa} \times I^{\epsilon_{\neg g}}_{\neg g\kappa}).$$

From the point of view of a person i of gender g, $I_{g\kappa}$ denotes the fraction of inventors of the same gender, whereas $I_{\neg g\kappa}$ denotes that for the opposite gender. The parameters ϵ_g and $\epsilon_{\neg g}$ capture the gender-specific relevance of role models while $\tau^E_{g\kappa}$ is the exposure distortion net of the role model frictions. Note that this formulation introduces a subtle yet significant externality within the model. Specifically, earlier generations of inventors might not recognize how their occupational decisions resonate and influence the career choices of subsequent cohorts. This market failure opens the door to affirmative action on the grounds of efficiency considerations, an issue I will revisit shortly.

In this extension of the model, the average research labor supply in each occupation is unchanged, but the proportion of individuals of gender g and cohort κ opting for research is now given by:

$$\mathrm{P}(z_i \geq \underline{z}_{g\kappa} \cap e_i = 1) = \underbrace{\frac{(1-\tau^E_{g\kappa})(1-\tau^L_{g\kappa})^\theta}{(1+\tau^H_{g\kappa})^{\frac{\theta\eta}{1+\nu}}}}_{\text{Distortions}} \times \underbrace{I^{\epsilon_g}_{g\kappa} I^{\epsilon_{\neg g}}_{\neg g\kappa}}_{\text{Role models}} \times \underbrace{\left(\frac{\omega^R_\kappa}{\omega^L_\kappa}\right)^\theta}_{\text{Wages}}$$

Hence, part of what was previously inferred as the exposure distortion is now captured by technological frictions from role models. To infer the residual exposure distortion, one must discipline the role model parameters ϵ_g and $\epsilon_{\neg g}$. To do so, I borrow estimates from Bell et al. (2018) who regress the fraction of children in a commuting zone who go on to patent in a specific technological field on the fraction of individuals of each gender from that commuting zone who were granted a patent in that same technological field.[19] In this extended model, the exposure distortion is subtly tempered relative to the baseline, and its interpretation becomes less straightforward. Nevertheless, conducting the counterfactual analysis in the environment, I find a slightly more modest increase in long-run income per person of 6.9% which translates to a consumption-equivalent welfare gain of 2.3%.

It was previously mentioned that this extension of the model introduces a novel externality: the current generation of individuals do not internalize that their occupation choice will influence the fraction of individuals in future cohorts who will have the opportunity to consider research as a career path. As such, the rationale behind subsidies to inventor wages goes beyond merely harnessing knowledge spillovers–it also serves as a corrective measure for this externality. This role model externality introduces the potential for transitional gender-specific optimal policy. Indeed, given an initially skewed gender composition, a welfare-maximizing planner might temporarily consider implementing *differential* wage subsidies for female inventors to expedite the transition towards a more efficient allocation of talent, trading off a slightly worse distribution of female inventive talent today in order to access a larger pool of talent in a nearer future.

Figure 11 charts the ratio of female to male inventor wage subsidies along a transition path. Initially, the planner optimally chooses to differentially subsidize the wages of female inventors by about 45%. Yet, this difference almost entirely tapers off within 50 years, eventually reaching gender parity. Hence, a planner seeking to dynamically maximize welfare would resort to affirmative action, at least temporarily, on the grounds of efficiency rather than equity. However, if such a policy were to be politically untenable, I find that the welfare cost of resorting to gender-neutral wage subsidies is negligible.

Technological Field Heterogeneity

Considering the nontrivial heterogeneity in the gender composition of technological fields presented in Figure 3(b), could inventive talent be misallocated not only between research and production but also across these different fields? To quantify this latter source of misallocation, the model can be extended to accommodate talent heterogeneity across technological fields.

In particular, the research sector could combine research labor R_{ft} from different

[19] The results of this regression are presented in Table 5 for both girls and boys. Bell et al. (2018) estimate statistically insignificant cross-gender coefficients but significant own-gender coefficients of 2.232 and 1.693 for girls and boys, respectively. To ensure coherence with the model, these estimates are first translated in elasticity form and then averaged to obtain parameter values of $\epsilon_g = 0.24$ and $\epsilon_{\neg g} = 0$.

Figure 11: Gender-Specific R&D Subsidy

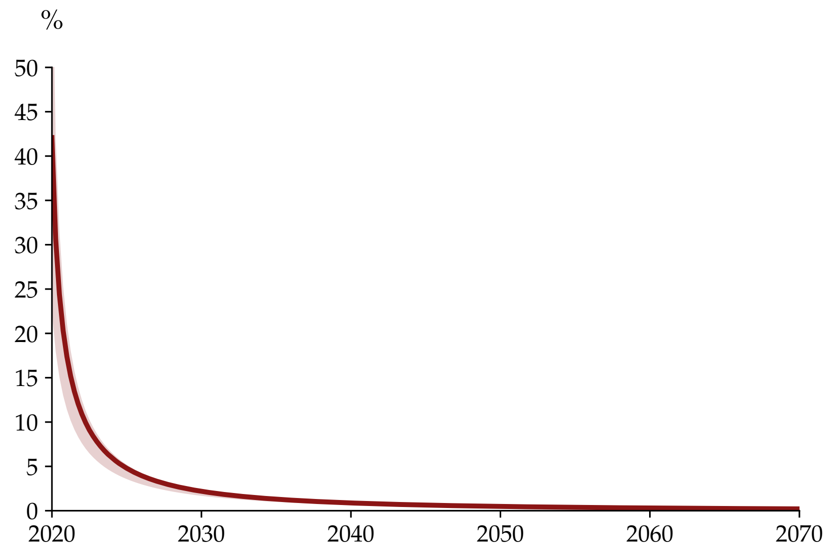

Note: This figure plots the ratio of optimal female to male inventor wage subsidies along a transition path. The shaded area corresponds to lower and upper bounds on the value of $\epsilon_g \in [0.14, 0.33]$ reflecting the 95% confidence interval around the estimates of Bell et al. (2018).

fields indexed by $f \in \{1, \ldots, F\}$ according to a Cobb-Douglas technology:

$$\dot{A}_t = A_t^{\phi} R_t \quad \text{where} \quad R_t = \prod_{f=1}^{F} R_{ft}^{\zeta_f} \quad \text{and} \quad \sum_{f=1}^{F} \zeta_f = 1$$

where the parameters $\boldsymbol{\zeta} \in [0,1]^F$ would govern the relative importance of each field in the production of new ideas. On the other side of this market would be individuals born with a vector of inventive talent $\mathbf{z}_i$ over those fields, drawn from a multivariate Pareto distribution with cumulative distribution function G:[20]

$$G(\mathbf{z}) = 1 - \left(\sum_{f=1}^{F} z_f^{\frac{-\theta}{1-\varrho}} \right)^{1-\varrho}.$$

The shape parameter $\theta > 1$ measures the degree of talent dispersion *across* individuals whereas $\varrho \in [0,1)$ determines their correlation between technological fields *within* individuals. In particular, talent draws are perfectly correlated across fields when $\varrho \to 1$,

[20]The talent scale parameters are normalized to unity for all fields as they play the same role as the parameters $\boldsymbol{\zeta}$. As described in Arkolakis, Rodríguez-Clare and Su (2017), this implies that the support of the distribution is $z_f \geq F^{1-\varrho}$ for all f.

or completely independent when $\varrho = 0$. With this field heterogeneity, exposure could be modeled as a vector of independent Bernoulli random variables with mean parameters $\{1 - \tau^E_{fg\kappa}\}_{f=1}^F \in [0,1]^F$.

A person's occupation choice would, therefore, be described in two successive steps. First, if they receive exposure to a subset of fields $\mathcal{F} \subseteq \{1, \ldots, F\}$, they identify the field f^* in that set that delivers the highest lifetime utility:

$$f^* = \arg\max_{f \in \mathcal{F}} \{U_{if}\}.$$

Then, they would decide whether to pursue that field or work in production, depending again on which of those two occupations promises the highest lifetime utility. Therefore, a person whose optimally chosen field is f would choose it over a career in production if and only if their talent in that field exceeds the field-specific threshold $\underline{z}_{fg\kappa}$:

$$\underline{z}_{fg\kappa} \equiv \frac{(1 + \tau^H_{fg\kappa})^\eta}{1 - \tau^L_{fg\kappa}} \times \frac{\omega^L_\kappa}{\omega^R_{f\kappa}}.$$

The following proposition characterizes the distribution of talent in optimally chosen fields conditional on exposure to a nonempty subset $\mathcal{F} \subseteq \{1, \ldots, F\}$:

Proposition 1. *Consider the people from cohort κ and gender g who received exposure to fields $\mathcal{F} \subseteq \{1, \ldots, F\}$. Solving and aggregating individual field choices for this group reveals that talent in chosen field $f \in \mathcal{F}$ follows a Pareto distribution:*[21]

$$P(\arg\max_{f' \in \mathcal{F}} \{U_{f'}\} = f \cap z_f \geq z) = \left(\frac{\mathcal{Z}^{\mathcal{F}}_{fg\kappa}}{z} \right)^\theta \quad \textit{where} \quad \mathcal{Z}^{\mathcal{F}}_{fg\kappa} \equiv \left(\frac{\underline{z}_{fg\kappa}^{\frac{\theta}{1-\varrho}}}{\sum_{f' \in \mathcal{F}} \underline{z}_{f'g\kappa}^{\frac{\theta}{1-\varrho}}} \right)^{\varrho/\theta}.$$

With this proposition, one can show that the total fraction $I_{g\kappa}$ of inventors in cohort κ and gender g is given by:

$$I_{g\kappa} \equiv \sum_{f=1}^{F} \sum_{\mathcal{F} \in \Gamma(f)} E^{\mathcal{F}}_{g\kappa} \left(\frac{\mathcal{Z}^{\mathcal{F}}_{fg\kappa}}{\underline{z}_{fg\kappa}} \right)^\theta \quad \text{where} \quad E^{\mathcal{F}}_{g\kappa} \equiv \prod_{f' \notin \mathcal{F}} \tau^E_{f'g\kappa} \times \prod_{f' \in \mathcal{F}} (1 - \tau^E_{f'g\kappa}).$$

Here, $\Gamma(f)$ denotes the set of all nonempty subsets of $\{1, \ldots, F\}$ containing element f and the $E^{\mathcal{F}}_{g\kappa}$ describes the probability that a person from cohort κ and gender g receives

[21]Proofs are in Appendix A.

exposure to the subset of fields $\mathcal{F}$:[22] One can similarly show that the average supply of research labor in technological field f, cohort κ, and gender g is proportional to:

$$\mathbb{E}[z_{if} \times h_i | z_{if} \geq \underline{z}_{fg\kappa}] \propto \frac{\omega^L_\kappa}{(1-\tau^L_{fg\kappa})\omega^R_{f\kappa}}.$$

Calibrating the parameters $\boldsymbol{\zeta}$ to match the distribution of inventors across fields and the parameter ϱ to target the cross-field inventive output correlation of multi-field inventors, I find that the long-run increase in income per person is almost unchanged from the baseline figure of 8.6%.[23] In fact, equalizing distortions across fields while leaving the aggregate share of women inventors unchanged barely moves income per person in the long run, suggesting that the bulk of talent misallocation is between the research and production sectors rather than across fields.

"On-The-Job" Human Capital

Women could be facing higher barriers to not only acquire, but also maintain/update human capital over their careers. For instance, it may prove more difficult for women to keep up with the technological frontier if the burden of childcare and housework falls disproportionately on them (Kim and Moser, 2021; Kaltenberg, Jaffe and Lachman, 2021). To reflect this possibility, I extend the theory to allow for gender-specific human capital depreciation: a reduced-form embodiment of gendered obstacles to skill maintenance. Under perfect foresight, it is straightforward to show that the differential depreciation rates are factored into consumption and saving decisions such that this distortion still operates through selection. That is, if gendered barriers to human capital maintenance constituted a prominent explanation for the scarcity of women in innovation, our theory implies that the latter would be more productive than their male colleagues at the onset of their careers, which is at odds with the evidence presented earlier.

[22] The set $\Gamma(f)$ has cardinality 2^{F-1} and summarizes every possible way in which a person can be exposed to field f.

[23] Specifically, I identify the set of multi-field inventors whose two main fields in terms of patents granted represent over a third of their lifetime inventive output. This set contains 118,061 inventors. From this bivariate distribution, I derive the empirical copula to nonparametrically estimate an upper tail dependence coefficient (UTDC) of 0.32 following the methodology of Frahm, Junker and Schmidt (2005). To relate this estimate to our model, we know from Arkolakis et al. (2017) that the UTDC of the multivariate Pareto distribution is equal to $2 - 2^{1-\varrho}$, which implies a value of $\varrho = 0.25$.

Occupational Preferences

Could gender differences in intrinsic preferences for innovation be a plausible candidate explanation for the underrepresentation of women among inventors? Broadening the model in this direction makes clear that this force yet again operates through selection. Indeed, if women disliked careers in innovation, only the most talented would pursue them despite their perceived disamenities. An alternative possibility is that individuals self-select into occupations based on those preferences rather than their inherent abilities or talents. If this is the case, our inference of the labor market distortion might not be accurate. As opposed to selection on talent, selection on preferences need not imply that female inventors would be more productive on average than their male counterparts.

Inventive Talent Uncertainty

How would our counterfactual calculations change if we were to relax the stylized assumption that individuals can perfectly observe their talent when choosing a career? If aspiring inventors face uncertainty about their own innovation potential, could gender differences in risk tolerance deter women from pursuing this occupation? To answer this question, I follow Bell et al. (2019) and assume that inventive talent x_i is the product of an observable *signal* z_i and an unobserved *shock* z_{iu} which is only realized after the occupation choice:

$$x_i = z_i \times z_{iu}.$$

The signal is still drawn from the distribution in equation (6), but the talent shock is instead drawn from a different Pareto distribution with a cumulative distribution function denoted by G_u:

$$G_u(z_u) = 1 - \left(\frac{\vartheta_u}{z_u}\right)^{\theta_u} \quad \text{where} \quad \vartheta_u \equiv \frac{\theta_u - 1}{\theta_u}.$$

Here, the shape parameter θ_u can be interpreted as the degree of talent uncertainty. With this formulation, the distribution's scale parameter ϑ_u directly depends on its shape parameter to preserve a unit average. In that sense, z_{iu} can be thought of as a mean-preserving spread to the distribution of the talent signal. This extension delivers a "fuzzy" selection into inventorship, as individuals with talent signals above the selection threshold may ultimately receive deceptively low talent shocks.

However, allowing for unobserved talent heterogeneity in our theory does not substantively affect the inference of distortions presented earlier. Indeed, if women were less tolerant of the uncertainty that inventive careers entail, only the most talented would earn enough to find it worthwhile to pursue this path. Since women and men inventors appear to be similarly productive, it seems unlikely that such gendered risk aversion could play a meaningful role. The "fuzzy" selection into inventorship that results from this uncertainty is also inconsequential for our counterfactual calculations. Letting θ_u take values of 1.1, 2, or infinity changes the long-run gain in income per person from eliminating distortions by less than a basis point.

6 Conclusion

Why are women so vastly and heterogeneously underrepresented among inventors? And what are the macroeconomic consequences of missing out on half of our brightest minds? To answer these questions, I develop a model of semi-endogenous growth in which individuals with heterogeneous inventive talent can choose between a career in innovation or production. However, three gendered barriers can deter or prevent women from pursuing their comparative advantage. They may face different forms of discrimination in the labor market, be confronted with higher obstacles to human capital formation, or lack the opportunities and role models to become innovators. Underlying this theory are two premises: (1) there are no intrinsic differences in inventive potential between women and men, and (2) exposure to inventive careers bears no relation to an individual's innate talent.

Interpreting micro-level data on U.S. inventors through the lens of this framework, I find that the underrepresentation of women is virtually all due to a lack of exposure to innovation. Women and men inventors are just too similarly productive, and the educational attainment gender gap among R&D workers is too narrow for distortions and frictions operating through selection or human capital to play a prominent role. From a policy perspective, this suggests that we ought to focus our attention and resources on bottlenecks that manifest earlier along the innovation pipeline.

I then take advantage of the general equilibrium structure of this theory to quantify the aggregate implications of lifting all barriers to female innovation. This calculation reveals that U.S. income per person would increase by 8.6% in the long run. Taking transition dynamics into account, eliminating all distortions would be equivalent to permanently raising everyone's consumption by 2.7%. Those economy-wide gains are mostly achieved by bringing better rather than more people into the process of innovation. Indeed, as barriers fall, new generations of ingenious women join the ranks of inventors, thus driving out marginally talented men.

This paper leaves the door open to many exciting avenues for future research. Which specific policies would be most effective in expanding access to inventive opportunities for young girls? Can data on inventor earnings, educational attainment, or childhood test scores provide more direct empirical evidence on various gendered distortions? Are people from low-income families and minority backgrounds facing the same obstacles to innovation as women? How much more prosperous would we be if we opened the doors of innovation to other underrepresented groups? Those are all outstanding but potentially fruitful questions that await future study.

References

Akcigit, Ufuk, Jeremy G Pearce, and Marta Prato, "Tapping into Talent: Coupling Education and Innovation Policies for Economic Growth," Working Paper 27862, National Bureau of Economic Research September 2020.

Alstott, Jeff, Ed Bullmore, and Dietmar Plenz, "powerlaw: A Python Package for Analysis of Heavy-Tailed Distributions," *PLoS ONE*, jan 2014, *9* (1), e85777.

American Chemical Society, "American Chemical Society National Historic Chemical Landmarks. Carl and Gerty Cori and Carbohydrate Metabolism.," http://www.acs.org/content/acs/en/education/whatischemistry/landmarks/carbohydratemetabolism.html (accessed July 1st, 2021) 2004.

— , "Gerty Theresa Cori (1896-1957)," https://www.acs.org/content/acs/en/education/whatischemistry/women-scientists/gerty-theresa-cori.html (accessed July 1st, 2021) 2004.

Arkolakis, Costas, Andrés Rodríguez-Clare, and Jiun-Hua Su, "A Multivariate Distribution with Pareto Tails and Pareto Maxima," Working Paper 2017.

— , Sun Kyoung Lee, and Michael Peters, "European Immigrants and the United States' Rise to the Technological Frontier," Working Paper 2020.

Atkeson, Andrew, Ariel T. Burstein, and Manolis Chatzikonstantinou, "Transitional Dynamics in Aggregate Models of Innovative Investment," *Annual Review of Economics*, 2019, *11* (1), 273–301.

Bell, Alex, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen, "Who Becomes an Inventor in America? The Importance of Exposure to Innovation," *The Quarterly Journal of Economics*, 11 2018, *134* (2), 647–713.

— , — , — , — , and — , "Joseph Schumpeter Lecture, EEA Annual Congress 2017: Do Tax Cuts Produce more Einsteins? The Impacts of Financial Incentives Versus Exposure to Innovation on the Supply of Inventors," *Journal of the European Economic Association*, 04 2019, *17* (3), 651–677.

Bento, Pedro, "Female Entrepreneurship in the US 1982-2012: Implications for Welfare and Aggregate Output," Working Paper 2021.

Blanchard, Olivier J., "Debt, Deficits, and Finite Horizons," *Journal of Political Economy*, 1985, *93* (2), 223–247.

Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb, "Are Ideas Getting Harder to Find?," *American Economic Review*, April 2020, *110* (4), 1104–44.

Bryan, Gharad and Melanie Morten, "The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia," *Journal of Political Economy*, 2019, *127* (5), 2229–2268.

Buera, Francisco J., Joseph P. Kaboski, and Yongseok Shin, "Finance and Development: A Tale of Two Sectors," *American Economic Review*, August 2011, *101* (5), 1964–2002.

Calvo, Guillermo A. and Maurice Obstfeld, "Optimal Time-Consistent Fiscal Policy with Finite Lifetimes," *Econometrica*, 1988, *56* (2), 411–432.

Card, David, "The Causal Effect of Education on Earnings," *Handbook of Labor Economics*, 1999, *3*, 1801–1863.

Carrell, Scott E., Marianne E. Page, and James E. West, "Sex and Science: How Professor Gender Perpetuates the Gender Gap," *The Quarterly Journal of Economics*, 08 2010, *125* (3), 1101–1144.

Celik, Murat Alp, "Does the Cream Always Rise to the Top? The Misallocation of Talent and Innovation," Working Paper 2022.

Chiplunkar, Gaurav and Pinelopi K Goldberg, "Aggregate Implications of Barriers to Female Entrepreneurship," Working Paper 28486, National Bureau of Economic Research February 2021.

Clauset, Aaron, Cosma Rohilla Shalizi, and M. E. J. Newman, "Power-Law Distributions in Empirical Data," *SIAM Review*, 2009, *51* (4), 661–703.

Einiö, Elias, Josh Feng, and Xavier Jaravel, "Social push and the direction of innovation," Working Paper 2022.

Ewens, Michael and Richard R. Townsend, "Are early stage investors biased against women?," *Journal of Financial Economics*, 2020, *135* (3), 653–677.

Frahm, Gabriel, Markus Junker, and Rafael Schmidt, "Estimating the tail-dependence coefficient: Properties and pitfalls," *Insurance: Mathematics and Economics*, 2005, *37* (1), 80–100. Papers presented at the DeMoSTAFI Conference, Québec, 20-22 May 2004.

Griliches, Zvi, "Patent Statistics as Economic Indicators: A Survey," *Journal of Economic Literature*, 1990, *28* (4), 1661–1707.

Hannon, Mary T, "The Patent Bar Gender Gap: Expanding the Eligibility Requirements to Foster Inclusion and Innovation in the US Patent System," *IP Theory*, 2021, *10*, 1.

Hochberg, Yael, Ali Kakhbod, Peiyao Li, and Kunal Sachdeva, "Inventor Gender and Patent Undercitation: Evidence from Causal Text Estimation," Working Paper 31592, National Bureau of Economic Research August 2023.

Hofstra, Bas, Vivek V. Kulkarni, Sebastian Munoz-Najar Galvez, Bryan He, Dan Jurafsky, and Daniel A. McFarland, "The Diversity & Innovation Paradox in Science," *Proceedings of the National Academy of Sciences*, 2020, *117* (17), 9284–9291.

Hsieh, Chang-Tai and Enrico Moretti, "Housing Constraints and Spatial Misallocation," *American Economic Journal: Macroeconomics*, April 2019, *11* (2), 1–39.

_, Erik Hurst, Charles I. Jones, and Peter J. Klenow, "The Allocation of Talent and U.S. Economic Growth," *Econometrica*, 2019, *87* (5), 1439–1474.

Hunt, Jennifer, Jean-Philippe Garant, Hannah Herman, and David J. Munroe, "Why are women underrepresented amongst patentees?," *Research Policy*, 2013, *42* (4), 831–843.

Jensen, Kyle, Balázs Kovács, and Olav Sorenson, "Gender Differences in Obtaining and Maintaining Patent Rights," *Nature biotechnology*, 2018, *36* (4), 307–309.

Jones, Charles I., "R & D-Based Models of Economic Growth," *Journal of Political Economy*, 1995, *103* (4), 759–784.

_, "Sources of U.S. Economic Growth in a World of Ideas," *American Economic Review*, March 2002, *92* (1), 220–239.

_, "The Past and Future of Economic Growth: A Semi-Endogenous Perspective," Working Paper 29126, National Bureau of Economic Research August 2021.

Kaltenberg, Mary, Adam B Jaffe, and Margie E Lachman, "Invention and the Life Course: Age Differences in Patenting," Working Paper 28769, National Bureau of Economic Research May 2021.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy, "Measuring Technological Innovation over the Long Run," *American Economic Review: Insights*, September 2021, *3* (3), 303–20.

Kim, Scott Daewon and Petra Moser, "Women in Science. Lessons from the Baby Boom," Working Paper 29436, National Bureau of Economic Research October 2021.

Kline, Patrick, Neviana Petkova, Heidi Williams, and Owen Zidar, "Who Profits from Patents? Rent-Sharing at Innovative Firms," *The Quarterly Journal of Economics*, 03 2019, *134* (3), 1343–1404.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, "Technological Innovation, Resource Allocation, and Growth," *The Quarterly Journal of Economics*, 03 2017, *132* (2), 665–712.

Lagakos, David and Michael E. Waugh, "Selection, Agriculture, and Cross-Country Productivity Differences," *American Economic Review*, April 2013, *103* (2), 948–80.

Lehr, Nils H., "R&D Return Dispersion and Economic Growth – The Case of Inventor Market Power," Working Paper 2023.

Morazzoni, Marta and Andrea Sy, "Female entrepreneurship, financial frictions and capital misallocation in the US," *Journal of Monetary Economics*, 2022, *129*, 93–118.

National Academies of Sciences, Engineering, and Medicine, *Sexual Harassment of Women: Climate, Culture, and Consequences in Academic Sciences, Engineering, and Medicine*, Washington, DC: The National Academies Press, 2018.

Pairolero, Nicholas, Andrew Toole, Charles DeGrazia, Mike Horia Teodorescu, and Peter-Anthony Pappas, "Closing the Gender Gap in Patenting: Evidence from a Randomized Control Trial at the USPTO," *Academy of Management Proceedings*, 2022, *2022* (1), 14401.

Peters, Michael, "Market Size and Spatial Growth - Evidence from Germany's Post-War Population Expulsions," Working Paper 29329, National Bureau of Economic Research October 2021.

Prato, Marta, "The Global Race for Talent: Brain Drain, Knowledge Transfer, and Economic Growth," Working Paper 2021.

Romer, Paul M., "Endogenous Technological Change," *Journal of Political Economy*, 1990, *98* (5, Part 2), S71–S102.

Ross, Matthew B., Britta M. Glennon, Raviv Murciano-Goroff, Enrico G. Berkes, Bruce A. Weinberg, and Julia I. Lane, "Women are Credited Less in Science than are Men," *Nature*, June 2022.

Roy, Andrew Donald, "Some Thoughts on the Distribution of Earnings," *Oxford Economic Papers*, 6 1951, *3* (2), 135–146.

Trimborn, Timo, Karl-Josef Koch, and Thomas M. Steger, "Multidimensional Transitional Dynamics: A Simple Numerical Procedure," *Macroeconomic Dynamics*, 2008, *12* (3), 301–319.

Yaari, Menahem E., "Uncertain Lifetime, Life Insurance, and the Theory of the Consumer," *The Review of Economic Studies*, 04 1965, *32* (2), 137–150.

A Theoretical Appendix

A.1 The Final Sector's Problem

Taking prices and the measure of varieties as given, the final sector's problem is to choose how much of each variety to produce with as to maximize profits. This is equivalent to the following cost minimization problem:

$$\min_{y_{jt}} \int_0^{A_t} p_{jt} y_{jt} \quad \text{s.t.} \quad \left(\int_0^{A_t} y_{jt}^{\frac{\sigma-1}{\sigma}} \mathrm{d}j \right)^{\frac{\sigma}{\sigma-1}} \geq Y_t.$$

The first-order conditions deliver the following demand functions:

$$y_{jt} = (\lambda_t / p_{jt})^{\sigma} Y_t$$

where $\lambda_t > 0$ is the Lagrange multiplier on the production constraint. This multiplier corresponds to the price of the final good P_t and can therefore be normalized to unity.

A.2 The Intermediate Sector's Problem

Taking the demand function for its variety, the rental rate of physical capital, and the wage paid to workers as given, the intermediate firm's problem is to choose its variety's price as well as physical capital and labor to maximize profits:

$$\pi_{jt} = \max_{p_{jt}, \ell_{jt}, k_{jt}} \{ p_{jt} y_{jt} - (r_t + \delta) k_{jt} - w_t^L \ell_{jt} \}.$$

In particular, the choice of physical capital and labor is the solution to the following cost minimization problem:

$$\min_{k_{jt}, \ell_{jt}} \{ (r_t + \delta) k_{jt} + w_t^L \ell_{jt} \} \quad \text{s.t.} \quad k_{jt}^{\alpha} \ell_{jt}^{1-\alpha} \geq y_{jt}.$$

The first-order conditions deliver the following demand functions:

$$k_{jt} = \frac{\lambda_t \alpha y_{jt}}{r_t + \delta} \quad \text{and} \quad \ell_{jt} = \frac{\lambda_t (1-\alpha) y_{jt}}{w_t^L}$$

where $\lambda_t > 0$ is the Lagrange multiplier on the production constraint. Substituting these demand functions in the firm's production function, one can solve for λ_t:

$$\lambda_t = \left(\frac{r_t + \delta}{\alpha}\right)^{\alpha} \left(\frac{w_t^L}{1-\alpha}\right)^{1-\alpha}$$

which is also equal to the firm's marginal cost of production. Substituting all the above demand functions in the firm's profit maximization problem, we obtain:

$$\pi_{jt} = \max_{p_{jt}} (p_{jt} - \lambda_t) Y_t / p_{jt}^{\sigma}.$$

The first-order condition delivers the following pricing function:

$$p_{jt} = \mu \left(\frac{r_t + \delta}{\alpha}\right)^{\alpha} \left(\frac{w_t^L}{1-\alpha}\right)^{1-\alpha} \quad \text{where} \quad \mu \equiv \frac{\sigma}{\sigma - 1}.$$

Notice that all firms make the same price, physical capital, and labor choices, which implies that profits are symmetric within the intermediate sector. Substituting the pricing function in the profit function and integrating across firms delivers:

$$\int_0^{A_t} \pi_{jt} \mathrm{d}j = \int_0^{A_t} p_{jt} y_{jt} \mathrm{d}j / \sigma = Y_t / \sigma.$$

Since profits are symmetric across firms, we have that:

$$\pi_{jt} = \frac{Y_t}{\sigma A_t} \quad \forall j \in [0, A_t].$$

Symmetry in physical capital and labor choices also implies that firms produce:

$$y_{jt} = \frac{K_t^{\alpha} L_t^{1-\alpha}}{A_t}.$$

Substituting this in the production function of the final good delivers:

$$Y_t = A_t^{\frac{1}{\sigma-1}} K_t^{\alpha} L_t^{1-\alpha}.$$

Similarly, by substituting the pricing function in the firm-level demand functions for physical capital and labor and integrating, we obtain the aggregate demand functions:

$$K_t = \frac{\alpha Y_t}{\mu (r_t + \delta)} \quad \text{and} \quad L_t = \frac{(1-\alpha) Y_t}{\mu w_t^L}$$

A.3 The Research Sector's Problem

Taking wages and the measure of varieties as given, the research sector's problem is to choose a patent price and research labor to maximize profits:

$$\max_{q_t, R_t} \{q_t A_t^{\phi} R_t - w_t^R R_t\}.$$

Hence, the first-order condition for research labor is:

$$q_t A_t^{\phi} = w_t^R.$$

Since there is free-entry in the intermediate sector, the final sector sets the price of a patent to extract all rents from the commercialization of an idea:

$$q_t = \int_t^{\infty} e^{-\int_t^{t'} r_\tau \mathrm{d}\tau} \pi_{t'} \mathrm{d}t'.$$

Differentiating with respect to time and using the expression for a firm's profits, we obtain the law of motion for the price of a patent:

$$\dot{q}_t = r_t q_t - \frac{Y_t}{\sigma A_t}.$$

Technological Field Heterogeneity

With heterogeneity across technological fields, taking wages and the measure of varieties as given, the research sector's problem is to choose a patent price and research labor of each type to maximize profits:

$$\max_{q_t, R_{ft}} \{q_t A_t^{\phi} \prod_{f=1}^{F} R_{ft}^{\zeta_f} - \sum_{f=1}^{F} w_t^R R_{ft}\}.$$

In particular, the choice of research labor in each field is the solution to the following cost-minimization problem:

$$\min_{R_{ft}} \sum_{f=1}^{F} w_{ft}^R R_{ft} \quad \text{s.t.} \quad \prod_{f=1}^{F} R_{ft}^{\zeta_f} \geq R_t.$$

The first-order conditions deliver the following demand functions:

$$R_{ft} = W_t \zeta_f R_t / w_{ft}^R$$

where $W_t > 0$ is the Lagrange multiplier on the aggregate research labor constraint. Substituting these demand functions in the expression for aggregate research labor, one can solve for W_t:

$$W_t = \prod_{f=1}^{F} \left(\frac{w_{ft}^R}{\zeta_f} \right)^{\zeta_f}$$

which is also equal to the research sector's marginal cost of production. Substituting the above demand functions in the research sector's profit maximization problem, we obtain:

$$\max_{q_t, R_t} \{ q_t A_t^{\phi} R_t - W_t R_t \}.$$

The first-order condition for aggregate research labor is:

$$q_t A_t^{\phi} = W_t$$

which, together with the demand functions for research labor of each type, delivers the corresponding market-clearing conditions:

$$w_{ft}^R = q_t A_t^{\phi} \zeta_f R_t / R_{ft}.$$

Finally, since there is free-entry in the intermediate sector, the final sector sets the price of a patent to extract all rents from the commercialization of an idea:

$$q_t = \int_t^{\infty} e^{-\int_t^{t'} r_\tau \mathrm{d}\tau} \pi_{t'} \mathrm{d}t'.$$

Differentiating with respect to time and using the expression for a firm's profits, we obtain the law of motion for the price of a patent:

$$\dot{q}_t = r_t q_t - \frac{Y_t}{\sigma A_t}.$$

A.4 The Individual's Problem

Taking prices as given, the problem of an individual i is to select a career in which they will choose consumption and schooling to maximize lifetime utility:

$$U_i = \max_{c_{it}, s_i} \int_{\kappa}^{\infty} e^{-(\rho+d)(t-\kappa)} \ln(c_{it}) \mathrm{d}t - (1 + \mathbb{1}_{\{i \in R\}} \tau_{g\kappa}^H) \beta s_i$$

subject to the flow budget constraint:

$$\dot{a}_{it} = \begin{cases} r_t a_{it} + (1-\tau^L_{gt}) w^R_t z_i h_i - c_{it} & \text{if } i \in R, \\ r_t a_{it} + w^L_t h_i - c_{it} & \text{if } i \in L, \end{cases}$$

and the initial condition $a_{i\kappa} = 0$. The corresponding current-value Hamiltonian is:

$$\mathcal{H}_t = \begin{cases} \ln(c_{it}) + \lambda_t [r_t a_{it} + (1-\tau^L_{gt}) w^R_t z_i h_i - c_{it}] & \text{if } i \in R, \\ \ln(c_{it}) + \lambda_t (r_t a_{it} + w^L_t h_i - c_{it}) & \text{if } i \in L \end{cases}$$

where λ_t denotes the costate variable and $\lim_{t\to\infty} e^{-(\rho+d)(t-\kappa)} \lambda_t a_{it} = 0$. The optimality conditions are:

$$\frac{\partial \mathcal{H}_t}{\partial c_{it}} = c_{it}^{-1} - \lambda_t = 0 \quad \text{and} \quad \frac{\partial \mathcal{H}_t}{\partial a_{it}} = \lambda_t r_t = (\rho+d)\lambda_t - \dot{\lambda}_t.$$

Combining those equations, we obtain the Euler equation and the No-Ponzi condition:

$$\frac{\dot{c}_{it}}{c_{it}} = r_t - \rho - d \quad \text{and} \quad \lim_{t\to\infty} e^{-\int_\kappa^t r_{t'} \mathrm{d}t'} a_{it} = 0.$$

Integrating the flow budget constraint using both equations delivers:

$$c_{it} = \begin{cases} (\rho+d)[a_{it} + (1-\tau^L_{gt}) \omega^R_t z_i h_i] & \text{if } i \in R, \\ (\rho+d)(a_{it} + \omega^L_t h_i) & \text{if } i \in L \end{cases} \quad \text{where} \quad \omega^o_t \equiv \int_t^\infty e^{-\int_t^{t'} r_\tau \mathrm{d}\tau} w^o_{t'} \mathrm{d}t'$$

for $o \in \{R, L\}$. Using the individual's Euler equation and the flow budget constraint's initial condition, we can express consumption in period t from the point of view of period κ as:

$$c_{it} = \begin{cases} (\rho+d)(1-\tau^L_{g\kappa}) \omega^R_\kappa z_i h_i e^{\int_\kappa^t (r_\tau - \rho - d)\mathrm{d}\tau} & \text{if } i \in R, \\ (\rho+d) \omega^L_\kappa h_i e^{\int_\kappa^t (r_\tau - \rho - d)\mathrm{d}\tau} & \text{if } i \in L. \end{cases}$$

Substituting this equation in the definition of lifetime utility:

$$U_i = \begin{cases} \frac{\ln[(\rho+d)(1-\tau^L_{g\kappa})\omega^R_\kappa z_i s_i^\eta] - 1}{\rho+d} - (1+\tau^H_{g\kappa})\beta s_i + \Delta_\kappa & \text{if } i \in R, \\ \frac{\ln[(\rho+d)\omega^L_\kappa s_i^\eta] - 1}{\rho+d} - \beta s_i + \Delta_\kappa & \text{if } i \in L, \end{cases}$$

where Δ_κ is defined for convenience as:

$$\Delta_\kappa \equiv \int_\kappa^\infty e^{-(\rho+d)(t-\kappa)} \int_\kappa^t r_{t'} \mathrm{d}t' \mathrm{d}t.$$

Choosing schooling time to maximize lifetime utility:

$$s_i = \frac{\eta}{\beta(\rho+d)(1+\mathbb{1}_{\{i\in R\}}\tau_{g\kappa}^H)}.$$

Substituting this choice back into the definition of lifetime utility:

$$U_i = \begin{cases} \frac{\ln[\eta^\eta(\rho+d)^{1-\eta}\hat{\omega}_{g\kappa}^R z_i] - \eta[\ln(\beta)+1]-1}{\rho+d} + \Delta_\kappa & \text{if } i \in R, \\ \frac{\ln[\eta^\eta(\rho+d)^{1-\eta}\omega_\kappa^L] - \eta[\ln(\beta)+1]-1}{\rho+d} + \Delta_\kappa & \text{if } i \in L, \end{cases}$$

where $\hat{\omega}_{g\kappa}^R$ is defined for convenience as:

$$\hat{\omega}_{g\kappa}^R \equiv \frac{1-\tau_{g\kappa}^L}{(1+\tau_{g\kappa}^H)^\eta} \times \omega_\kappa^R.$$

Individual i will decide to pursue research if and only if their lifetime utility in that career exceeds that in production, which delivers a selection threshold on talent:

$$\underline{z}_{g\kappa} \equiv \frac{\omega_\kappa^L}{\hat{\omega}_{g\kappa}^R}.$$

Technological Field Heterogeneity

With heterogeneity across technological fields, taking prices as given, the problem of an individual i is to select a career in which they will choose consumption and schooling to maximize lifetime utility:

$$U_i = \max_{c_{it}, s_i} \int_\kappa^\infty e^{-(\rho+d)(t-\kappa)} \ln(c_{it}) \mathrm{d}t - (1+\mathbb{1}_{\{i\in f\}}\tau_{fg\kappa}^H)\beta s_i$$

subject to the flow budget constraint:

$$\dot{a}_{it} = \begin{cases} r_t a_{it} + (1-\tau_{fgt}^L) w_{ft}^R z_{if} h_i - c_{it} & \text{if } i \in f, \\ r_t a_{it} + w_t^L h_i - c_{it} & \text{if } i \in L, \end{cases}$$

and the initial condition $a_{i\kappa} = 0$. The corresponding current-value Hamiltonian is:

$$\mathcal{H}_t = \begin{cases} \ln(c_{it}) + \lambda_t[r_t a_{it} + (1-\tau^L_{fgt})w^R_{ft} z_{if} h_i - c_{it}] & \text{if } i \in f, \\ \ln(c_{it}) + \lambda_t(r_t a_{it} + w^L_t h_i - c_{it}) & \text{if } i \in L \end{cases}$$

where λ_t denotes the costate variable and $\lim_{t\to\infty} e^{-(\rho+d)(t-\kappa)}\lambda_t a_{it} = 0$. The optimality conditions are:

$$\frac{\partial \mathcal{H}_t}{\partial c_{it}} = c_{it}^{-1} - \lambda_t = 0 \quad \text{and} \quad \frac{\partial \mathcal{H}_t}{\partial a_{it}} = \lambda_t r_t = (\rho+d)\lambda_t - \dot{\lambda}_t.$$

Combining those equations, we obtain the Euler equation and the No-Ponzi condition:

$$\frac{\dot{c}_{it}}{c_{it}} = r_t - \rho - d \quad \text{and} \quad \lim_{t\to\infty} e^{-\int_\kappa^t r_{t'}\mathrm{d}t'} a_{it} = 0.$$

Integrating the flow budget constraint using both equations delivers:

$$c_{it} = \begin{cases} (\rho+d)[a_{it} + (1-\tau^L_{fgt})\omega^R_{ft} z_{if} h_i] & \text{if } i \in f, \\ (\rho+d)(a_{it} + \omega^L_t h_i) & \text{if } i \in L \end{cases}$$

where ω^L_t and ω^R_{ft} are defined as before. Using the individual's Euler equation and the flow budget constraint's initial condition, we can express consumption in period t from the point of view of period κ as:

$$c_{it} = \begin{cases} (\rho+d)(1-\tau^L_{fg\kappa})\omega^R_{f\kappa} z_{if} h_i e^{\int_\kappa^t (r_\tau - \rho - d)\mathrm{d}\tau} & \text{if } i \in f, \\ (\rho+d)\omega^L_\kappa h_i e^{\int_\kappa^t (r_\tau - \rho - d)\mathrm{d}\tau} & \text{if } i \in L. \end{cases}$$

Substituting this equation in the definition of lifetime utility:

$$U_i = \begin{cases} \frac{\ln[(\rho+d)(1-\tau^L_{fg\kappa})\omega^R_{f\kappa} z_{if} s_i^\eta] - 1}{\rho+d} - (1+\tau^H_{fg\kappa})\beta s_i + \Delta_\kappa & \text{if } i \in f, \\ \frac{\ln[(\rho+d)\omega^L_\kappa s_i^\eta] - 1}{\rho+d} - \beta s_i + \Delta_\kappa & \text{if } i \in L. \end{cases}$$

Choosing schooling time to maximize lifetime utility:

$$s_i = \frac{\eta}{\beta(\rho+d)(1+\mathbb{1}_{\{i\in f\}}\tau^H_{fg\kappa})}.$$

Substituting this choice back into the definition of lifetime utility:

$$U_i = \begin{cases} \frac{\ln[\eta^\eta(\rho+d)^{1-\eta}\hat{\omega}^R_{fg\kappa} z_{if}] - \eta[\ln(\beta)+1]-1}{\rho+d} + \Delta_\kappa & \text{if } i \in f, \\ \frac{\ln[\eta^\eta(\rho+d)^{1-\eta}\omega^L_\kappa] - \eta[\ln(\beta)+1]-1}{\rho+d} + \Delta_\kappa & \text{if } i \in L, \end{cases}$$

where $\hat{\omega}^R_{fg\kappa}$ is defined for convenience as:

$$\hat{\omega}^R_{fg\kappa} \equiv \frac{1-\tau^L_{fg\kappa}}{(1+\tau^H_{fg\kappa})^\eta} \times \omega^R_{f\kappa}.$$

Individual i will decide to pursue research in their optimally chosen field f if and only if their lifetime utility in that career exceeds that in production, which delivers a selection threshold on talent:

$$\underline{z}_{fg\kappa} \equiv \frac{\omega^L_\kappa}{\hat{\omega}^R_{fg\kappa}}.$$

A.5 Aggregation

The fraction of individuals from cohort κ in period t is equal to $be^{-b(t-\kappa)}$. Denoting average consumption within cohort κ as $c_t(\kappa)$ delivers the following expression for average consumption in the economy:

$$c_t = \int_{-\infty}^{t} be^{-b(t-\kappa)} c_t(\kappa) \mathrm{d}\kappa.$$

Differentiating with respect to time and using the Euler equation:

$$\dot{c}_t = (r_t - \rho - d)c_t - b[c_t - c_t(t)].$$

Substituting in the consumption function and the asset market clearing condition:

$$\dot{c}_t = (r_t - \rho - d)c_t - b(\rho + d)[k_t + q_t A_t / N_t + H_t - H_t(t)]$$

where $k_t \equiv K_t/N_t$ is physical capital per person, and H_t and $H_t(t)$ denote average human wealth across all cohorts and in the most recent one, respectively:[24]

$$H_t \equiv \frac{\omega_t^L L_t}{N_t} + \frac{\sum_g (1-\tau_{gt}^L)\omega_t^R R_{gt}}{2N_t} \quad \text{and} \quad H_t(t) \equiv \hat{\eta}[1+(\vartheta-1)I_t]\omega_t^L.$$

Here, the fraction I_t of inventors among the most recent cohort is defined as:

$$I_t \equiv \frac{1}{2} \times \sum_g (1-\tau_{gt}^E)\underline{z}_{gt}^{-\theta}.$$

The laws of motion for L_t and R_{gt} are:

$$\dot{L}_t = b\hat{\eta}(1-I_t)N_t - dL_t \quad \text{and} \quad \dot{R}_{gt} = \frac{b\hat{\eta}\vartheta\omega_t^L I_{gt} N_t}{(1-\tau_{gt}^L)\omega_t^R} - dR_{gt}$$

and the laws of motion for Δ_t and ω_t^R are:

$$\dot{\Delta}_t = (\rho+d)\Delta_t - \frac{r_t}{\rho+d} \quad \text{and} \quad \dot{\omega}_t^R = r_t\omega_t^R - w_t^R.$$

Define the normalized variable $x_t^* \equiv x_t e^{-g_x t}$ such that g_x is the balanced growth rate of variable x_t and let $N_t^* = 1$ for all t. Then, collecting the above and performing some simple substitutions, we obtain the system of ordinary differential equations describing the dynamics of the equilibrium allocation:

$$\begin{aligned}
\dot{c}_t^* &= (r_t - \rho - d - g_c)c_t^* - b(\rho+d)[k_t^* + q_t^* A_t^* + H_t^* - H_t^*(t)], \\
\dot{L}_t^* &= b[\hat{\eta}(1-I_t) - L_t^*], \\
\dot{R}_{gt}^* &= b\{(\hat{\eta}\vartheta\omega_t^{L*} I_{gt})/[(1-\tau_{gt}^L)\omega_t^{R*}] - R_{gt}^*\}, \\
\dot{\omega}_t^* &= (r_t - g_c)\omega_t^{R*} - w_t^{R*}, \\
\dot{A}_t^* &= A_t^{*\phi}(\textstyle\sum_g R_{gt}^*/2) - g_A A_t^*, \\
\dot{q}_t^* &= (r_t - g_q)q_t^* - y_t^*/(\sigma A_t^*), \\
\dot{k}_t^* &= y_t^* - (\delta + n + g_c)k_t^* - c_t^*, \\
\dot{\Delta}_t &= (\rho+d)\Delta_t - r_t/(\rho+d)
\end{aligned}$$

[24]The two constants ϑ and $\hat{\eta}$ are defined as $\vartheta \equiv \theta/(\theta-1)$ and $\hat{\eta} \equiv \{\eta/[\beta(\rho+d)]\}^{\eta}$.

where we have the additional definitions:

$$g_c \equiv n/[(1-\alpha)(\sigma-1)(1-\phi)],$$
$$g_A \equiv n/(1-\phi),$$
$$g_q \equiv n + g_c - g_A.$$

The steady state of the model is found by setting all time derivatives to zero and solving the resulting nonlinear system of equations.

Technological Field Heterogeneity

Proof of Proposition 1. Consider the group of people from cohort κ and gender g who received exposure to fields $\mathcal{F} \subseteq \{1, \ldots, F\}$ and choose field f:

$$U_{f'} \leq U_f \quad \Leftrightarrow \quad z_{f'} \leq \frac{\hat{\omega}_{fg\kappa}}{\hat{\omega}_{f'g\kappa}} \times z_f \quad \forall f' \in \mathcal{F} \backslash f.$$

Using property (ii) of the multivariate Pareto distribution in Arkolakis et al. (2017), we have that the lower-dimensional marginal of G for the nonempty subset $\mathcal{F}$ is:

$$G^{\mathcal{F}}(\mathbf{z}) = 1 - \left(\sum_{f' \in \mathcal{F}} z_{f'}{}^{\frac{-\theta}{1-\varrho}} \right)^{1-\varrho}.$$

Without loss of generality, suppose that $\mathcal{F}$ has cardinality $n \in \{1, \ldots, F\}$ such that we can order its elements from $\mathcal{F}_1$ to $\mathcal{F}_n$. Then, we know that:

$$\mathrm{P}(\arg\max_{f' \in \mathcal{F}}\{U_{f'}\} = f \cap z_f = z) = \frac{\partial G^{\mathcal{F}}\left(z_{\mathcal{F}_1} \leq \frac{\hat{\omega}_{fg\kappa}}{\hat{\omega}_{\mathcal{F}_1 g\kappa}} \times z, \ldots, z_f \leq z, \ldots, z_{\mathcal{F}_n} \leq \frac{\hat{\omega}_{fg\kappa}}{\hat{\omega}_{\mathcal{F}_n g\kappa}} \times z\right)}{\partial z_f}.$$

Using the definition of $G^{\mathcal{F}}$, we can rewrite the above equation as:

$$\mathrm{P}(\arg\max_{f' \in \mathcal{F}}\{U_{f'}\} = f \cap z_f = z) = \theta \mathcal{Z}_{fg\kappa}^{\mathcal{F}}{}^{\theta} z^{-\theta-1}$$

where the term $\mathcal{Z}_{fg\kappa}^{\mathcal{F}}$ is defined as:

$$\mathcal{Z}_{fg\kappa}^{\mathcal{F}} \equiv \left(\frac{\underline{z}_{fg\kappa}^{\frac{\theta}{1-\varrho}}}{\sum_{f' \in \mathcal{F}} \underline{z}_{f'g\kappa}^{\frac{\theta}{1-\varrho}}} \right)^{\varrho/\theta}.$$

$\square$

With heterogeneity across technological fields, we have:[25]

$$H_t \equiv \frac{\omega_t^L L_t}{N_t} + \frac{\sum_{f=1}^{F} \sum_g (1-\tau_{fgt}^L)\omega_{ft}^R R_{fgt}}{2N_t} \quad \text{and} \quad H_t(t) \equiv \hat{\eta}[1+(\vartheta-1)I_t]\omega_t^L.$$

Here, the fraction I_t of inventors among the most recent cohort is defined as:

$$I_t \equiv \frac{1}{2} \times \sum_g \sum_{f=1}^{F} \sum_{\mathcal{F}\in\Gamma(f)} E_{gt}^{\mathcal{F}} \left(\frac{\mathcal{Z}_{fgt}^{\mathcal{F}}}{\underline{z}_{fgt}} \right)^{\theta}.$$

The laws of motion for L_t and R_{fgt} are:

$$\dot{L}_t = b\hat{\eta}(1-I_t)N_t - dL_t \quad \text{and} \quad \dot{R}_{fgt} = \frac{b\hat{\eta}\vartheta\omega_t^L I_{fgt} N_t}{(1-\tau_{fgt}^L)\omega_{ft}^R} - dR_{fgt} \quad \forall f \in \{1,\dots,F\}$$

and the laws of motion for Δ_t and ω_{ft}^R are:

$$\dot{\Delta}_t = (\rho+d)\Delta_t - \frac{r_t}{\rho+d} \quad \text{and} \quad \dot{\omega}_{ft}^R = r_t\omega_{ft}^R - w_{ft}^R \quad \forall f \in \{0,\dots,F\}.$$

Collecting the above and performing some simple substitutions, we obtain the system of ordinary differential equations describing the dynamics of the equilibrium allocation:

$$\begin{aligned}
\dot{c}_t^* &= (r_t - \rho - d - g_c)c_t^* - b(\rho+d)[k_t^* + q_t^* A_t^* + H_t^* - H_t^*(t)], \\
\dot{L}_t^* &= b[\hat{\eta}(1-I_t) - L_t^*], \\
\dot{R}_{fgt}^* &= b\{(\hat{\eta}\vartheta\omega_t^{L*} I_{fgt})/[(1-\tau_{fgt}^L)\omega_{ft}^{R*}] - R_{fgt}^*\}, \\
\dot{\omega}_{ft}^* &= (r_t - g_c)\omega_{ft}^{R*} - w_t^{R*}, \\
\dot{A}_t^* &= A_t^{*\phi} \textstyle\prod_{f=1}^{F} (\sum_g R_{fgt}^*/2)^{\zeta_f} - g_A A_t^*, \\
\dot{q}_t^* &= (r_t - g_q)q_t^* - y_t^*/(\sigma A_t^*), \\
\dot{k}_t^* &= y_t^* - (\delta + n + g_c)k_t^* - c_t^*, \\
\dot{\Delta}_t &= (\rho+d)\Delta_t - r_t/(\rho+d).
\end{aligned}$$

[25]The two constants ϑ and $\hat{\eta}$ are defined as $\vartheta \equiv \theta/(\theta-1)$ and $\hat{\eta} \equiv \{\eta/[\beta(\rho+d)]\}^{\eta}$.

A.6 Social Welfare

Define the utilitarian social welfare function as:

$$
\begin{aligned}
W_t = & \int_{-\infty}^{t} be^{-b(t-\kappa)} \int_{t}^{\infty} e^{-(\rho+d)(\tau-t)} \mathbb{E}[\ln(c_{i\tau})] \mathrm{d}\tau \mathrm{d}\kappa \\
& + \int_{t}^{\infty} be^{-(\rho-n)(\kappa-t)} \int_{\kappa}^{\infty} e^{-(\rho+d)(\tau-\kappa)} \mathbb{E}[\ln(c_{i\tau})] \mathrm{d}\tau \mathrm{d}\kappa
\end{aligned}
$$

where the first and second terms, respectively correspond to the remaining lifetime utility of surviving and future cohorts. For simplicity, the term involving preferences over pre-career leisure is ignored as all individuals receive the same disutility from schooling, regardless of distortions. Changing the order of integration:

$$
W_t = \int_{t}^{\infty} e^{-(\rho-n)(\tau-t)} \int_{-\infty}^{\tau} be^{-b(\tau-\kappa)} \mathbb{E}[\ln(c_{i\tau})] \mathrm{d}\kappa \mathrm{d}\tau.
$$

Balanced growth path.—Suppose the economy was and remains on its balanced growth path. Then, average flow utility for cohort κ at time τ is:

$$
\mathbb{E}[\ln(c_{i\tau})] = \ln[\hat{\eta}(\rho+d)\omega_{\tau}^{L}] + (r-\rho-d-g_c)(\tau-\kappa) + I/\theta.
$$

Integrating across cohorts, we have:

$$
\int_{-\infty}^{\tau} be^{-b(\tau-\kappa)} \mathbb{E}[\ln(c_{i\tau})] \mathrm{d}\kappa = \ln[\hat{\eta}(\rho+d)\omega_{\tau}^{L}] + (r-\rho-d-g_c)/b + I/\theta.
$$

Further integrating over time, we obtain:

$$
W_t^{BGP} = \{\ln[\hat{\eta}(\rho+d)\omega_t^{L*}] + (r-\rho-d-g_c)/b + I/\theta + g_c/(\rho-n)\}/(\rho-n).
$$

Transition path.—Suppose instead that the economy was on its balanced growth path before time t and is launched on a transition path thereafter. Then, average flow utility for *future* cohort κ at time τ is:

$$
\mathbb{E}[\ln(c_{i\tau})|\kappa \geq t] = \ln[\hat{\eta}(\rho+d)\omega_{\kappa}^{L}] + \int_{\kappa}^{\tau} (r'_t - \rho - d) \mathrm{d}t' + I_{\kappa}/\theta.
$$

Integrating over time, we have:

$$
\int_{\kappa}^{\infty} e^{-(\rho+d)(\tau-\kappa)} \mathbb{E}[\ln(c_{i\tau})|\kappa \geq t] \mathrm{d}\tau = \{\ln[\hat{\eta}(\rho+d)\omega_{\kappa}^{L}] + I_{\kappa}/\theta - 1\}/(\rho+d) + \Delta_{\kappa}.
$$

Further integrating over future cohorts, we obtain:

$$W_t^F = \frac{b\int_t^\infty e^{-(\rho-n)(\kappa-t)}\{\ln[\hat{\eta}(\rho+d)\omega_\kappa^{L*}] + I_\kappa/\theta + (\rho+d)\Delta_\kappa\}\mathrm{d}\kappa}{\rho+d} + \frac{b(g_c+n-\rho)}{(\rho+d)(\rho-n)^2}.$$

In contrast, average flow utility for *surviving* cohort κ at time τ is:

$$\begin{aligned}\mathbb{E}[\ln(c_{i\tau})|\kappa<t] = \ln[\hat{\eta}(\rho+d)\omega_t^{L^{\text{old}}}] + \int_t^\tau (r'_t-\rho-d)\mathrm{d}t' \\ + (1-I)\ln\left[e^{(r-\rho-d-g_c)(t-\kappa)} + \frac{\omega_t^{L^{\text{new}}}}{\omega_t^{L^{\text{old}}}} - 1\right] \\ + \sum_g I_g\left\{\ln\left[e^{(r-\rho-d-g_c)(t-\kappa)} + \frac{(1-\tau_{gt}^{L,\text{new}})\omega_t^{R^{\text{new}}}}{(1-\tau_g^{L,\text{old}})\omega_t^{R^{\text{old}}}} - 1\right] + \frac{1}{\theta}\right\}\end{aligned}$$

where the superscript "new" indexes a variable on the new transition path and the superscript "old" indexes a variable on the previous balanced growth path. Integrating over time, we have:

$$\begin{aligned}&\int_t^\infty e^{-(\rho+d)(\tau-t)}\mathbb{E}[\ln(c_{i\tau})|\kappa<t]\mathrm{d}\tau = \{\ln[\hat{\eta}(\rho+d)\omega_t^{L^{\text{old}}}] - 1\}/(\rho+d) + \Delta_t \\ &+ (1-I)\ln\left[e^{(r-\rho-d-g_c)(t-\kappa)} + \frac{\omega_t^{L^{\text{new}}}}{\omega_t^{L^{\text{old}}}} - 1\right]/(\rho+d) \\ &+ \sum_g I_g\left\{\ln\left[e^{(r-\rho-d-g_c)(t-\kappa)} + \frac{(1-\tau_{gt}^{L,\text{new}})\omega_t^{R^{\text{new}}}}{(1-\tau_g^{L,\text{old}})\omega_t^{R^{\text{old}}}} - 1\right] + \frac{1}{\theta}\right\}/(\rho+d).\end{aligned}$$

Further integrating over surviving cohorts, we obtain:

$$\begin{aligned}&W_t^S = \{\ln[\hat{\eta}(\rho+d)\omega_t^{L^{\text{old}*}}] - 1\}/(\rho+d) + \Delta_t \\ &+ (1-I)\int_{-\infty}^t be^{-b(t-\kappa)}\ln\left[e^{(r-\rho-d-g_c)(t-\kappa)} + \frac{\omega_t^{L^{\text{new}*}}}{\omega_t^{L^{\text{old}*}}} - 1\right]\mathrm{d}\kappa/(\rho+d) \\ &+ \sum_g I_g\left\{\int_{-\infty}^t be^{-b(t-\kappa)}\ln\left[e^{(r-\rho-d-g_c)(t-\kappa)} + \frac{(1-\tau_{gt}^{L,\text{new}})\omega_t^{R^{\text{new}*}}}{(1-\tau_g^{L,\text{old}})\omega_t^{R^{\text{old}*}}} - 1\right]\mathrm{d}\kappa + \frac{1}{\theta}\right\}/(\rho+d).\end{aligned}$$

Consumption-equivalent welfare.—By how much should we permanently increase the consumption of all individuals from the balanced growth path economy to leave them as well off as the individuals from the transition path economy? Letting the proportional permanent increase in consumption be denoted by λ, the answer to that question satisfies:

$$\lambda = \exp[(\rho - n)(W_t^S + W_t^F - W_t^{BGP})].$$

A.7 Theoretical Extensions

In this section of the Appendix, I consider several extensions to the theoretical framework presented in the main text.

Role Models and Affirmative Action

Introducing role models in this framework is straightforward. The only difference relative to the aggregation results above is that the fraction $I_{gt}(t)$ of inventors among the most recent cohort is defined as:

$$I_{gt}(t) \equiv (1 - \tau_{gt}^E)\underline{z}_{gt}^{-\theta} I_{gt}^{\epsilon_g} I_{\neg gt}^{\epsilon_{\neg g}}$$

where I_{gt} and I_t are now defined as:

$$I_{gt} \equiv \int_{-\infty}^{t} be^{-b(t-\kappa)} I_t(\kappa) \mathrm{d}\kappa \quad \text{and} \quad I_t \equiv \frac{1}{2} \times \sum_g I_{gt}.$$

Differentiating this equation with respect to time delivers:

$$\dot{I}_{gt} = b[I_{gt}(t) - I_{gt}]$$

which consists of two additional ordinary differential equations to include in the system.

Let us now derive a constrained optimal allocation of resources under this theoretical extension. In particular, assume that a benevolent social planner seeks to maximize the present discounted value of consumption per capita while respecting privately optimal individual schooling decisions. Then, the planner's stationary current-value

Hamiltonian is given by:

$$
\begin{aligned}
\mathcal{H}_t = \ln(c_t^*) &+ \lambda_t^k [A_t^{*\frac{1}{\sigma-1}} k_t^{*\alpha} L_t^{*1-\alpha} - (\delta + n + g_c) k_t^* - c_t^*] \\
&+ \lambda_t^L b[\hat{\eta}(1 - \sum_g \underline{z}_{gt}^{-\theta} I_{gt}^{\epsilon_g} I_{\neg gt}^{\epsilon_{\neg g}}/2) - L_t^*] \\
&+ \lambda_t^A (A_t^{*\phi} \sum_g R_{gt}^*/2 - g_A A_t^*) \\
&+ b \sum_g \lambda_{gt}^R (\hat{\eta} \vartheta \underline{z}_{gt}^{1-\theta} I_{gt}^{\epsilon_g} I_{\neg gt}^{\epsilon_{\neg g}} - R_{gt}^*) \\
&+ b \sum_g \lambda_{gt}^I (\underline{z}_{gt}^{-\theta} I_{gt}^{\epsilon_g} I_{\neg gt}^{\epsilon_{\neg g}} - I_{gt}).
\end{aligned}
$$

Assuming as in our parameterization that $\epsilon_g = \epsilon$ and $\epsilon_{\neg g} = 0$, the optimality conditions are given by:

$$
\begin{aligned}
\frac{\partial \mathcal{H}_t}{\partial c_t^*} &= 1/c_t^* - \lambda_t^k = 0 \\
\frac{\partial \mathcal{H}_t}{\partial \underline{z}_{gt}} &= \lambda_t^L b \hat{\eta} \theta \underline{z}_{gt}^{-\theta-1} I_{gt}^{\epsilon}/2 - \lambda_{gt}^R \theta b \hat{\eta} \underline{z}_{gt}^{-\theta} I_{gt}^{\epsilon} - \lambda_{gt}^I \theta b \underline{z}_{gt}^{-\theta-1} I_{gt}^{\epsilon} = 0, \\
\frac{\partial \mathcal{H}_t}{\partial k_t^*} &= \lambda_t^k [\alpha A_t^{*\frac{1}{\sigma-1}} (L_t^*/k_t^*)^{1-\alpha} - \delta - n - g_c] = (\rho - n)\lambda_t^k - \dot{\lambda}_t^k, \\
\frac{\partial \mathcal{H}_t}{\partial L_t^*} &= \lambda_t^k (1-\alpha) A_t^{*\frac{1}{\sigma-1}} (k_t^*/L_t^*)^{\alpha} - \lambda_t^L b = (\rho - n)\lambda_t^L - \dot{\lambda}_t^L, \\
\frac{\partial \mathcal{H}_t}{\partial A_t^*} &= \lambda_t^k A_t^{*\frac{2-\sigma}{\sigma-1}} k_t^{*\alpha} L_t^{*1-\alpha}/(\sigma-1) + \lambda_t^A (\phi A_t^{*\phi-1} \sum_g R_{gt}/2 - g_A) = (\rho - n)\lambda_t^A - \dot{\lambda}_t^A, \\
\frac{\partial \mathcal{H}_t}{\partial R_{gt}^*} &= \lambda_t^A A_t^{\phi}/2 - \lambda_{gt}^R b = (\rho - n)\lambda_{gt}^R - \dot{\lambda}_{gt}^R, \\
\frac{\partial \mathcal{H}_t}{\partial I_{gt}} &= -\lambda_t^L b \hat{\eta} \epsilon \underline{z}_{gt}^{-\theta} I_{gt}^{\epsilon-1}/2 + \lambda_{gt}^R b \hat{\eta} \vartheta \epsilon \underline{z}_{gt}^{1-\theta} I_{gt}^{\epsilon-1} + \lambda_{gt}^I b(\epsilon \underline{z}_{gt}^{-\theta} I_{gt}^{\epsilon-1} - 1) = (\rho - n)\lambda_{gt}^I - \dot{\lambda}_{gt}^I.
\end{aligned}
$$

Rearranging these equations, we obtain:

$$
\begin{aligned}
c_t^* &= 1/\lambda_t^k \\
\underline{z}_{gt} &= (\lambda_t^L/2 - \lambda_{gt}^I/\hat{\eta})/\lambda_{gt}^R, \\
\dot{\lambda}_t^k &= [\rho + \delta + g_c - \alpha A_t^{*\frac{1}{\sigma-1}} (L_t^*/k_t^*)^{1-\alpha}]\lambda_t^k, \\
\dot{\lambda}_t^L &= (\rho + d)\lambda_t^L - (1-\alpha) A_t^{*\frac{1}{\sigma-1}} (k_t^*/L_t^*)^{\alpha} \lambda_t^k, \\
\dot{\lambda}_t^A &= (\rho - n + g_A - \phi A_t^{*\phi-1} \sum_g R_{gt}^*/2)\lambda_t^A - A_t^{*\frac{2-\sigma}{\sigma-1}} k_t^{*\alpha} L_t^{*1-\alpha} \lambda_t^k/(\sigma-1), \\
\dot{\lambda}_{gt}^R &= (\rho + d)\lambda_{gt}^R - \lambda_t^A A_t^{*\phi}/2, \\
\dot{\lambda}_{gt}^I &= (\rho + d - b\epsilon \underline{z}_{gt}^{-\theta} I_{gt}^{\epsilon-1})\lambda_{gt}^I + b\hat{\eta}\epsilon \underline{z}_{gt}^{-\theta} I_{gt}^{\epsilon-1} \lambda_t^L/2 - b\hat{\eta}\vartheta\epsilon \underline{z}_{gt}^{1-\theta} I_{gt}^{\epsilon-1} \lambda_{gt}^R.
\end{aligned}
$$

If the planner was instead constrained to choose a single talent threshold $\underline{z}_t$ for all g, we would obtain:

$$
\begin{aligned}
c_t^* &= 1/\lambda_t^k \\
\underline{z}_t &= (\lambda_t^L \textstyle\sum_g I_{gt}^\epsilon/2 - \sum_g \lambda_{gt}^I I_{gt}^\epsilon/\hat{\eta})/\sum_g \lambda_{gt}^R I_{gt}^\epsilon, \\
\dot{\lambda}_t^k &= [\rho + \delta + g_c - \alpha A_t^{*\frac{1}{\sigma-1}}(L_t^*/k_t^*)^{1-\alpha}]\lambda_t^k, \\
\dot{\lambda}_t^L &= (\rho + d)\lambda_t^L - (1-\alpha)A_t^{*\frac{1}{\sigma-1}}(k_t^*/L_t^*)^\alpha \lambda_t^k, \\
\dot{\lambda}_t^A &= (\rho - n + g_A - \phi A_t^{*\phi-1} \textstyle\sum_g R_{gt}^*/2)\lambda_t^A - A_t^{*\frac{2-\sigma}{\sigma-1}} k_t^{*\alpha} L_t^{*1-\alpha}\lambda_t^k/(\sigma-1), \\
\dot{\lambda}_{gt}^R &= (\rho + d)\lambda_{gt}^R - \lambda_t^A A_t^{*\phi}/2, \\
\dot{\lambda}_{gt}^I &= (\rho + d - b\epsilon \underline{z}_t^{-\theta} I_{gt}^{\epsilon-1})\lambda_{gt}^I + b\hat{\eta}\epsilon \underline{z}_t^{-\theta} I_{gt}^{\epsilon-1}\lambda_t^L/2 - b\hat{\eta}\vartheta\epsilon \underline{z}_t^{1-\theta} I_{gt}^{\epsilon-1}\lambda_{gt}^R.
\end{aligned}
$$

Gendered Preferences

Suppose that lifetime utility took the form:

$$U_i = \int_\kappa^\infty e^{-(\rho+d)(t-\kappa)} \ln(c_{it}) \mathrm{d}t + \mathbb{1}_{\{i \in f\}} \ln(\nu_{fg}) - \beta s_i$$

where $\nu_{fg} > 0$ measures gendered preferences for particular technological fields. Since this extension does not affect the choices of consumption or schooling, the selection threshold on talent becomes:

$$\underline{z}_{fg\kappa} \equiv \frac{(1+\tau_{g\kappa}^H)^\eta}{1-\tau_{g\kappa}^L} \times \frac{\omega_\kappa^L}{\omega_{f\kappa}^R} \times \frac{1}{\nu_{fg}}.$$

Therefore, lower preferences for specific fields raise the talent threshold above which it becomes worthwhile to pursue those fields. In other words, if women disliked being inventors, only the most talented would earn enough to compensate for a field's disamenity. In particular, this threshold implies that average research productivity in field f, gender g, and cohort κ is proportional to:

$$\mathbb{E}[z_{if} \times h_i | z_{if} \geq \underline{z}_{fg\kappa}] \propto \frac{1}{1-\tau_{g\kappa}^L} \times \frac{1}{\nu_{fg}}.$$

Therefore, if inherited gendered preferences constituted a plausible explanation for the scarcity of women in innovation, the latter would have to be more productive than their male colleagues, above and beyond what would be implied by the labor market distortion alone.

"On-The-Job" Human Capital

Suppose that over the course of their research career, the human capital of a person of gender g evolves as follows:

$$\dot{h}_{it} = -\mathbb{1}_{\{i \in R\}} \delta_g \times h_{it} \quad \text{where} \quad h_{i\kappa} = s_i^{\eta}.$$

Solving the linear differential equation delivers:

$$h_{it} = s_i^{\eta} e^{-\delta_g (t-\kappa)}.$$

Then, the individual consumption function becomes:

$$c_{it} = \begin{cases} (\rho + d)(1 - \tau_{g\kappa}^L)\omega_{g\kappa}^R z_i s_i^{\eta} e^{\int_{\kappa}^{t} (r_{\tau} - \rho - d) \mathrm{d}\tau} & \text{if } i \in R, \\ (\rho + d)\omega_{\kappa}^L s_i^{\eta} e^{\int_{\kappa}^{t} (r_{\tau} - \rho - d) \mathrm{d}\tau} & \text{if } i \in L, \end{cases}$$

where the stream of future wages in research is now defined as:

$$\omega_{g\kappa}^R \equiv \int_{\kappa}^{\infty} e^{-\int_{\kappa}^{t'} (r_{\tau} + \delta_g) \mathrm{d}\tau} w_{t'}^R \mathrm{d}t'.$$

Since this extension does not affect the choices of consumption and schooling, the selection threshold on talent becomes:

$$\underline{z}_{g\kappa} \equiv \frac{(1 + \tau_{g\kappa}^H)^{\eta}}{1 - \tau_{g\kappa}^L} \times \frac{\omega_{\kappa}^L}{\omega_{g\kappa}^R}.$$

On a balanced growth path, where the rental rate of capital is constant, we can rewrite:

$$\underline{z}_{gt} \equiv \frac{(1 + \tau_{fg}^H)^{\eta}}{1 - \tau_{fg}^L} \times \frac{w_t^{L*}}{w_t^{R*}} \times \frac{r + \delta_g - g_c}{r - g_c}.$$

Therefore, a larger rate of human capital depreciation raises the talent threshold above which it becomes worthwhile to pursue innovation. This threshold implies that average research productivity among gender g and cohort t is proportional to:

$$\mathbb{E}[z_i \times h_i | z_i \geq \underline{z}_{gt}] \propto \frac{r + \delta_g - g_c}{1 - \tau_g^L}.$$

That is, if women inventors faced higher barriers to maintaining/updating their knowledge and skills on-the-job, they would tend to be more productive than their male colleagues at the onset of their career, above and beyond what would be implied by the labor market distortion alone.

Inventive Talent Uncertainty and Gendered Risk Aversion

Suppose that expected lifetime utility took the form:

$$U_i = \mathbb{E} \int_{\kappa}^{\infty} e^{-(\rho+d)(t-\kappa)} u_g(c_{it}) \mathrm{d}t - \beta s_i e^{\psi_g (t-\kappa)} \quad \text{where} \quad u_g(c) = \frac{c^{1-\gamma_g} - 1}{1-\gamma_g}$$

and where $\gamma_g > 1$ denotes the coefficient of relative risk aversion of gender group g, and ψ_g should be assumed to be chosen such that schooling time is stationary on a balanced growth path. Suppose further that inventive talent x_i is the product of a heterogeneous observable signal z_i and an unobserved shock z_{iu} which is only realized after the occupation choice:

$$x_i = z_i \times z_{iu}$$

where the talent shock is drawn from a univariate Pareto distribution with a cumulative distribution function denoted by G_u:

$$G_u(z_u) = 1 - \left(\frac{\vartheta_u}{z_u} \right)^{\theta_u} \quad \text{where} \quad \vartheta_u \equiv \frac{\theta_u - 1}{\theta_u}.$$

In this case, the current-value Hamiltonian of the individual's problem is:

$$\mathcal{H}_t = \begin{cases} u_g(c_{it}) + \lambda_t [r_t a_{it} + (1-\tau_{gt}^L) w_t^R x_i h_i - c_{it}] & \text{if } i \in R, \\ u_g(c_{it}) + \lambda_t (r_t a_{it} + w_t^L h_i - c_{it}) & \text{if } i \in L, \end{cases}$$

where λ_t denotes the costate variable and $\lim_{t\to\infty} e^{-(\rho+d)(t-\kappa)} \lambda_t a_{it} = 0$. The optimality conditions are:

$$\frac{\partial \mathcal{H}_t}{\partial c_{it}} = c_{it}^{-\gamma_g} - \lambda_t = 0 \quad \text{and} \quad \frac{\partial \mathcal{H}_t}{\partial a_{it}} = \lambda_t r_t = (\rho+d)\lambda_t - \dot{\lambda}_t.$$

Combining those equations, we obtain the Euler equation and the No-Ponzi condition:

$$\frac{\dot{c}_{it}}{c_{it}} = (r_t - \rho - d)/\gamma_g \quad \text{and} \quad \lim_{t\to\infty} e^{-\int_{\kappa}^{t} r_{t'} \mathrm{d}t'} a_{it} = 0.$$

Integrating the flow budget constraint using both equations delivers:

$$c_{it} = \begin{cases} [a_{it} + (1 - \tau^L_{gt})\omega^R_t x_i h_i]/\Delta_{gt} & \text{if } i \in R, \\ (a_{it} + \omega^L_t h_i)/\Delta_{gt} & \text{if } i \in L, \end{cases}$$

where we have the following two definitions:

$$\omega^o_t \equiv \int_t^\infty e^{-\int_t^{t'} r_\tau \mathrm{d}\tau} w^o_{t'} \mathrm{d}t' \quad \text{and} \quad \Delta_{gt} \equiv \int_t^\infty e^{-\int_t^{t'} [(\gamma_g - 1)r_\tau + \rho + d]\mathrm{d}\tau/\gamma_g} \mathrm{d}t'$$

for $o \in \{R, L\}$. Using the individual's Euler equation and the flow budget constraint's initial condition, we can express consumption in period t from the point of view of period κ as:

$$c_{it} = \begin{cases} (1 - \tau^L_{g\kappa})\omega^R_\kappa x_i h_i e^{\int_\kappa^t (r_\tau - \rho - d)\mathrm{d}\tau/\gamma_g}/\Delta_{g\kappa} & \text{if } i \in R, \\ \omega^L_\kappa h_i e^{\int_\kappa^t (r_\tau - \rho - d)\mathrm{d}\tau/\gamma_g}/\Delta_{g\kappa} & \text{if } i \in L. \end{cases}$$

In terms of timing, let us assume that individuals first choose which career path they want to pursue, after which the talent shock is realized and becomes actionable information to guide consumption and education decisions. Therefore, substituting the above equation in the definition of lifetime utility delivers:

$$U_i = \begin{cases} \frac{[(1-\tau^L_{g\kappa})\omega^R_\kappa x_i h_i]^{1-\gamma_g}\Delta_{g\kappa}^{\gamma_g} - 1}{1-\gamma_g} - (1 + \tau^H_{g\kappa})\beta s_i e^{\psi_g(t-\kappa)} & \text{if } i \in R, \\ \frac{(\omega^L_\kappa h_i)^{1-\gamma_g}\Delta_{g\kappa}^{\gamma_g} - 1}{1-\gamma_g} - \beta s_i e^{\psi_g(t-\kappa)} & \text{if } i \in L. \end{cases}$$

Choosing schooling time to maximize lifetime utility:

$$s_i = \begin{cases} \left\{ \frac{\eta[(1-\tau^L_{g\kappa})\omega^R_\kappa x_i]^{1-\gamma_g}\Delta_{g\kappa}^{\gamma_g}}{(1+\tau^H_{g\kappa})\beta e^{\psi_g(t-\kappa)}} \right\}^{\frac{1}{1+\eta(\gamma_g-1)}} & \text{if } i \in R, \\ \left\{ \frac{\eta {\omega^L_\kappa}^{1-\gamma_g}\Delta_{g\kappa}^{\gamma_g}}{\beta e^{\psi_g(t-\kappa)}} \right\}^{\frac{1}{1+\eta(\gamma_g-1)}} & \text{if } i \in L. \end{cases}$$

Substituting this choice back into the definition of *expected* lifetime utility:

$$U_i = \begin{cases} \frac{[1+\eta(\gamma_g-1)]\mathbb{E}\left\{[(\eta^\eta \hat{\omega}^R_{g\kappa} z_i z_{iu} e^{-\eta\psi_g(t-\kappa)}/\beta^\eta)^{1-\gamma_g}\Delta_{g\kappa}^{\gamma_g}]^{\frac{1}{1+\eta(\gamma_g-1)}}\right\} - 1}{1-\gamma_g} & \text{if } i \in R, \\ \frac{[1+\eta(\gamma_g-1)][(\eta^\eta \omega^L_\kappa e^{-\eta\psi_g(t-\kappa)}/\beta^\eta)^{1-\gamma_g}\Delta_{g\kappa}^{\gamma_g}]^{\frac{1}{1+\eta(\gamma_g-1)}} - 1}{1-\gamma_g} & \text{if } i \in L, \end{cases}$$

where $\hat{\omega}_{g\kappa}$ is defined as usual. Individual i will decide to pursue research if and only if their talent signal is above the threshold:

$$\underline{z}_{g\kappa} \equiv \frac{\omega_\kappa^L}{\hat{\omega}_{g\kappa}^R} \times \frac{\theta_u}{\theta_u - 1} \times \left\{ \frac{[1+\eta(\gamma_g - 1)]\theta_u}{[1+\eta(\gamma_g-1)]\theta_u + \gamma_g - 1} \right\}^{\frac{1+\eta(\gamma_g-1)}{\gamma_g - 1}}.$$

The additional last term reflects how talent uncertainty influences the occupation choice and is decreasing in θ_u, but increasing in γ_g and η. Intuitively, a larger degree of talent uncertainty and risk aversion, as captured by a low value of θ_u and a high value of γ_g, discourages those who receive relatively low talent signals and for whom taking the risk of drawing yet another low talent shock is simply not worth it.

The role of η is slightly more subtle. As mentioned earlier, the timing of decisions is such that talent uncertainty is resolved after the occupation choice, but *before* the schooling choice. With human capital and talent being complements, schooling decisions amplify the degree of productivity dispersion ex-post and, therefore, the degree of earnings uncertainty ex-ante. Consequently, the easier it is to turn schooling into productive human capital, as measured by a larger value for η, the riskier the inventor occupation seems. To illustrate these points, notice that in a risk-neutral world without human capital (γ_g and η equal to zero), this talent uncertainty term would simply be equal to one.

Now, consider the people from cohort κ and gender g whose talent signal exceeds the selection threshold $\underline{z}_{g\kappa}$. Denoting by $\hat{g}$ the corresponding probability density function of talent signals, the distribution of total inventive talent x for this group must satisfy:

$$\mathrm{P}(x_i = x | z_i \geq \underline{z}_{g\kappa}) = \int_{-\infty}^{\infty} \hat{g}(x/z_u) g_u(z_u) |z_u|^{-1} \mathrm{d}z_u.$$

Denote that probability density function by $g(x)$. Since the respective supports of $\hat{g}$ and g_u are $(\underline{z}_{g\kappa}, +\infty)$ and $(\vartheta_u, +\infty)$, the integrand in the above expression is nonzero if and only if $\vartheta_u < z_u < x/\underline{z}_{g\kappa}$. With the definitions of $\hat{g}$ and g_u, this implies:

$$g(x) = \begin{cases} \theta^2 \left(\frac{\vartheta_u \underline{z}_{g\kappa}}{x}\right)^\theta \ln\left(\frac{x}{\vartheta_u \underline{z}_{g\kappa}}\right) / x & \text{if } \theta_u = \theta, \\ \left(\frac{\theta_u \theta}{\theta_u - \theta}\right) \left[\left(\frac{\vartheta_u \underline{z}_{g\kappa}}{x}\right)^\theta - \left(\frac{\vartheta_u \underline{z}_{g\kappa}}{x}\right)^{\theta_u} \right] / x & \text{otherwise} \end{cases}$$

with support $x > \vartheta_u \times \underline{z}_{g\kappa}$. To have a sense of what this distribution looks like, Figure 12 plots its probability density function for $\theta = 2$ and different values of θ_u. The right tail of that distribution still follows a power law with tail exponent $\min\{\theta, \theta_u\}$, but its left tail is hump-shaped, in contrast to the initial Pareto distributions. This hump-shape pattern is a consequence of the "fuzzy" selection into inventorship, as individuals with talent signals above the selection threshold may ultimately receive deceptively low talent shocks.

Figure 12: The Distribution of Inventive Talent with Uncertainty

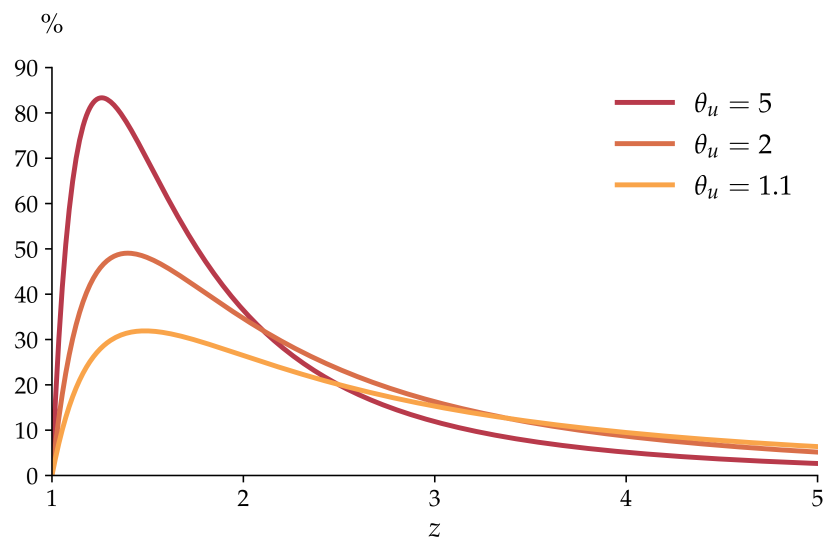

Note: This figure plots $g(x)$ for different values of θ_u when $\theta = 2$. For lower values of θ_u, there is more dispersion in talent shocks, which delivers a thicker right tail of inventive talent.

Taking the product of talent and human capital and integrating over the resulting distribution delivers an expression for average research productivity among gender g and cohort κ:

$$\mathbb{E}[x_i \times h_i | z_i \geq \underline{z}_{g\kappa}] \propto \frac{1}{1 - \tau^L_{g\kappa}} \times \left\{ \frac{[1 + \eta(\gamma_g - 1)]\theta_u}{[1 + \eta(\gamma_g - 1)]\theta_u + \gamma_g - 1} \right\}^{\frac{1}{\gamma_g - 1}} .$$

Note that average research productivity is still inversely proportional to the "keep rate" of the labor market tax but is also an increasing function of γ_g. Therefore, if larger relative risk aversion is what was keeping women away from innovation, we would expect women inventors to be, once again, more productive than their male colleagues, above and beyond what would be implied by the labor market tax alone.