

Who Profits from Patents?

Rent-Sharing at Innovative Firms*

Patrick Kline
UC-Berkeley

Neviana Petkova
US Department of Treasury

Heidi Williams
MIT

Owen Zidar
Princeton

November 2, 2018

Abstract

This paper analyzes how patent-induced shocks to labor productivity propagate into worker compensation using a new linkage of US patent applications to US business and worker tax records. We infer the causal effects of patent allowances by comparing firms whose patent applications were initially allowed to those whose patent applications were initially rejected. To identify patents that are ex-ante valuable, we extrapolate the excess stock return estimates of Kogan et al. (2017) to the full set of accepted and rejected patent applications based on predetermined firm and patent application characteristics. An initial allowance of an ex-ante valuable patent generates substantial increases in firm productivity and worker compensation. By contrast, initial allowances of lower ex-ante value patents yield no detectable effects on firm outcomes. Patent allowances lead firms to increase employment, but entry wages and workforce composition are insensitive to patent decisions. On average, workers capture roughly 30 cents of every dollar of patent-induced surplus in higher earnings. This share is roughly twice as high among workers present since the year of application. These earnings effects are concentrated among men and workers in the top half of the earnings distribution, and are paired with corresponding improvements in worker retention among these groups. We interpret these earnings responses as reflecting the capture of economic rents by senior workers, who are most costly for innovative firms to replace.

*We are very grateful to Joey Anderson, Ivan Badinski, Jeremy Brown, Alex Fahey, Sam Grondahl, Stephanie Kestelman, Tamri Matishvili, Mahnum Shahzad, Karthik Srinivasan, and John Wieselthier for excellent research assistance, and to Daron Acemoglu, Lawrence Katz, Bentley MacLeod, John Van Reenen, four anonymous referees and seminar participants at Brown, Dartmouth, the LSE/IFS/STICERD seminar, Michigan, Michigan State, MIT, NBER Labor Studies, NBER Productivity, Northwestern, NYU, Princeton, Stanford, the Toulouse Network for Information Technology, UC-Berkeley, UC-Irvine, UIUC, Chicago Booth, and the Washington Center for Equitable Growth for helpful comments. The construction of some of the data analyzed in this publication was supported by the National Institute on Aging and the NIH Common Fund, Office of the NIH Director, through Grant U01-AG046708 to the National Bureau of Economic Research (NBER); the content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH or NBER. This work/research was also funded by the Ewing Marion Kauffman Foundation; the contents of this publication are solely the responsibility of the Grantee. Financial support from the Alfred P. Sloan Foundation, NSF Grant Numbers 1151497 and 1752431, the Washington Center for Equitable Growth, the Kathryn and Grant Swick Faculty Research Fund, and the Booth School of Business at the University of Chicago are also gratefully acknowledged. Opinions expressed in the paper do not necessarily represent the views and policies of the US Department of the Treasury. Corresponding author: Williams; MIT Department of Economics; 50 Memorial Drive; E52-440; Cambridge MA 02142; heidiw@mit.edu.

1 Introduction

Competitive models of labor markets are predicated on the notion that firms have no power to set wages. However, there is mounting empirical evidence that firms contribute substantially to wage inequality among identically skilled workers (Card, Heining, and Kline 2013; Card, Cardoso, and Kline 2016; Barth et al. 2016; Jäger 2016; Goldschmidt and Schmieder 2017; Helpman et al. 2017; Abowd, McKinney, and Zhao 2018; Sorkin 2018; Song et al. forthcoming). This emerging evidence has renewed interest in mechanisms through which variation in firm productivity can influence worker pay (see Lentz and Mortensen 2010; Manning 2011 for reviews).

While a sizable empirical literature has documented that fluctuations in firm performance and worker compensation are strongly related (Card et al., 2018), these correlations are open to widely varying interpretations. Early studies (e.g., Christofides and Oswald 1992; Blanchflower, Oswald, and Sanfey 1996) estimated industry-level relationships that could simply reflect competitive market dynamics. A second generation of studies (Van Reenen 1996; Hildreth 1998; Abowd, Kramarz, and Margolis 1999) used firm-level data to study how shocks to firm performance translate into worker pay, but was unable to adjust for potential changes in worker composition. More recent work (Guiso, Pistaferri, and Schivardi 2005; Card, Devicienti, and Maida 2014; Friedrich et al. 2014; Card, Cardoso, and Kline 2016; Carlsson, Messina, and Skans 2016; Lamadon 2016; Mogstad et al. 2017) adjusts for composition biases by examining the comovement between changes in firm productivity and the wage growth of incumbent workers. However, observational fluctuations in standard labor productivity measures are likely to reflect a number of factors (e.g., market-wide fluctuations in product demand, changes in non-pecuniary firm amenities, or drift in labor market institutions) that can influence wages without necessarily signaling a violation of price-taking behavior by firms.

In this paper, we analyze how patent-induced shocks to firm performance propagate into worker compensation. Patent allowances offer a useful source of variation because they provide firms with well-defined monopoly rights that can yield a prolonged stream of potentially substantial economic rents. Standard models of frictional labor markets (e.g., Pissarides 2000; Hall and Milgrom 2008; Pissarides 2009) suggest that these product market rents will be shared with workers whenever the employment relationship is (re-)negotiated, yet surprisingly little is known about how broadly patent-generated rents are shared in practice.

Our analysis relies on a new linkage of two datasets: (i) the census of published patent applications submitted to the US Patent and Trademark Office (USPTO) between roughly 2001 and 2011 and (ii) the universe of US Treasury business tax filings and worker earnings histories drawn from W2 and 1099 tax filings. The business tax filings data offers a uniquely high-quality set of firm-level variables, from which we are able to construct multiple measures of firm performance. Likewise, the business and worker tax filings provide a unique window into compensation outcomes for many different types of workers, including firm officers and owners, who prevail

at the top of the income distribution (Smith et al., 2017).

We infer the causal effect of patent allowances by comparing firms whose applications were initially allowed to those whose applications were initially rejected. Within so-called “art units” (technological areas designated by the USPTO), firms with initially allowed and initially rejected applications submitted in the same year are found to exhibit similar levels and trends in outcomes prior to their initial patent decision. We also document that initial patent decisions are difficult to predict based on firm characteristics or geography, corroborating the view that these decisions constitute truly idiosyncratic — as opposed to market level — shocks.

It is well-known that most patents generate little ex-post value to the firm (Pakes 1986; Hall, Jaffe, and Trajtenberg 2001). We build on insights from two recent studies to identify a subsample of valuable patents that induce meaningful shifts in firm outcomes at the time the patents are allowed. First, following the work of Farre-Mensa, Hegde, and Ljungqvist (2017), we restrict our analysis to firms applying for a patent for the first time, for which patent decisions are likely to be more consequential. Second, among this sample of first-time applicants, we build on the analysis of Kogan et al. (2017) who use event studies to estimate the excess stock market return realized on the grant date of US patents assigned to publicly traded firms. Specifically, we develop a methodology for extrapolating Kogan et al.’s patent value estimates to the non-publicly traded firms in our sample, and to firms whose patent applications are never granted. We use characteristics of firms and their patent applications that are fixed at the time of application as the basis for extrapolating patent values, and show that these value estimates are strong predictors of treatment effect heterogeneity in our sample. These value estimates also provide us with an additional validation of our research design: patents with low predicted value are found to have economically small and statistically insignificant effects on firm performance and worker compensation.

Using these data, we then investigate the consequences of obtaining an ex-ante valuable patent allowance for firm performance and worker compensation, and relate our findings to different explanations for the propagation of firm-specific shocks into worker wages. Corroborating recent research based on US Census data (Balasubramanian and Sivadasan 2011), we find that firm size and average labor productivity rise rapidly in response to initial allowances of ex-ante valuable patents. The average wage and salary income of workers at these firms rises in tandem with measures of average labor productivity. An allowance of a patent application in the top quintile of ex-ante predicted value raises firm-level surplus — defined as the sum of W2 earnings and business earnings before interest, taxes, and depreciation — by roughly \$12,400 per W2 employee per year, while W2 earnings at the firm rise by approximately \$3,700 per worker per year.

Patent allowances not only raise average earnings at assignee firms, but also exacerbate within-firm inequality on a variety of margins. Earnings impacts are heavily concentrated among employees in the top quartile of the within-firm earnings distribution and among employees listed on firm tax returns as “firm officers.” Likewise, we find that the earnings of owner-operators rise more than those of other employees. Earnings of male employees rise

strongly in response to a patent allowance, while the earnings of female employees are less responsive to patent decisions.

A handful of previous studies have investigated how inventor wages change in response to patent applications or patent grants (Toivanen and Väänänen 2012; Depalo and Di Addario 2014; Bell et al. 2016; Aghion et al. 2018). Consistent with these results, we find that the earnings of “inventors” — defined as employees ever listed as inventors on a patent application as in Bell et al. (2016) — respond to patent allowance decisions. Inventor earnings are more responsive to patent allowance decisions than are the earnings of non-inventors, similar to the findings presented in contemporaneous work by Aghion et al. (2018), which analyzes how inventor and non-inventor earnings in Finnish firms evolve before and after patent applications are filed.

While these impacts on firm aggregates could, in principle, be confounded by compositional changes, we find no evidence that innovative firms upgrade the quality of their workforce in response to patent allowances. Although patent allowances lead firms to expand by hiring slightly younger workers, the average prior earnings of both new hires and firm separators is unaffected by patent decisions, suggesting that there are no major changes in the skill composition of worker inflows to or outflows from the firm on a year-to-year basis.

Different theoretical frameworks offer divergent predictions about how firm-specific shocks will affect the wages of new and incumbent workers. Empirically, the earnings of workers who were employed by the firm in the year of application respond very strongly to patent decisions. Having a valuable patent allowed raises the average earnings of these “firm stayers” by roughly \$7,800 — or approximately 11% — per year. These gains appear to be concentrated among firm stayers who, in the year of application, were located in the top half of the firm’s earnings distribution. We also find that the earnings of male firm stayers respond more strongly to patent allowances than those of female firm stayers, which are estimated to be positive, albeit somewhat imprecise. By contrast, we are unable to detect any response of entry wages to patent allowances, which is inconsistent with the predictions of both static wage posting models and traditional bargaining models involving Nash-style surplus splitting at the time of hiring (Pissarides 2000; Hall and Milgrom 2008; Pissarides 2009). While some dynamic wage posting models (e.g., Postel-Vinay and Robin 2002) can generate drops in entry wages in response to a productivity increase, these models predict greater wage *growth* for new hires, a phenomenon for which we also find no evidence. A candidate explanation for such “insider/outsider” distinctions in earnings impacts is that the wage fluctuations of incumbent workers represent changes in market perceptions of a worker’s underlying ability (Gibbons and Murphy 1992; Holmström 1999). However, we find much smaller and statistically insignificant earnings effects for workers who leave the firm, suggesting that our results are unlikely to be driven by public learning about worker quality.

To interpret our findings, we sketch a simple model in which incumbent workers are imperfectly substitutable with new hires. As in Becker (1964), Stevens (1994), and Manning (2006), this mechanism provides an avenue for incumbents to extract rents from the firm in the form of wage premia. Motivated by this framework, we fit a series of

“rent-sharing” specifications analogous to standard cost-price pass-through specifications used to study imperfect competition in product markets (Goldberg and Hellerstein, 2013; Weyl and Fabinger, 2013; Gorodnichenko and Talavera, 2017). Using patent decisions as an instrument for firm surplus, we find that worker earnings rise by roughly 29 cents of every dollar of patent allowance-induced surplus, with an approximate elasticity of 0.35, which is comparable to the earlier estimates of Abowd and Lemieux (1993) and Van Reenen (1996) that were based on firm-level aggregates. Importantly, failing to instrument for surplus yields smaller elasticities, closer to those in the recent studies reviewed by Card et al. (2018) that assume statistical innovations to average labor productivity constitute structural productivity shocks. Consistent with our model, rent-sharing with firm stayers is more pronounced than it is with average workers: stayers capture roughly 61 cents of every dollar of surplus for an approximate elasticity of 0.56. When we exclude employees ever listed as inventors on a patent, pass-through to firm stayers falls to roughly 48 cents with a corresponding elasticity of 0.5. Though this elasticity estimate is larger than what has been found by most previous studies, its 90% confidence interval encompasses many estimates in the literature.

In our model, firms share rents with incumbent workers to increase the odds of retaining them. We provide event study evidence that retention rises in response to patent allowances, with larger responses among workers in the top half of the earnings distribution. The fact that groups experiencing the largest earnings responses exhibit the largest retention responses strongly suggests that the earnings fluctuations we measure constitute rents, rather than, say, risk-sharing arrangements that hold workers to a participation constraint (Holmström 1979, 1989). Using the patent decision as an instrument for wages, we estimate a retention-wage elasticity of roughly 1.2, with a 90% confidence interval ranging from 0.46 to 3.1. When converted to a separation-wage elasticity, our point estimate lies near the middle of the range of quasi-experimental estimates reviewed in Manning (2011).

Viewed through the lens of our model, our point estimates imply that incumbent workers capture roughly 73% of their replacement costs in wage premia. We also estimate that the marginal replacement cost of an incumbent worker at a firm receiving a patent allowance is roughly equal to a new hire’s annual earnings. These findings suggest that separations of key personnel can be extremely costly to innovative firms, even when these employees are not themselves inventors. More broadly, our results suggest that the influence of firm conditions on worker wages depends critically on their degree of replaceability, which may be influenced both by the duration of the relationship between the worker and firm and by a worker’s position within the firm hierarchy, issues emphasized in recent empirical studies of wage setting at European firms by Buhai et al. (2014), Jäger (2016), and Garin and Silverio (2017). In contrast with European settings, however, the legal barriers to hiring and firing workers are comparatively minimal for the set of newly innovative US firms which are the focus of our analysis. The fact that seniority appears to mediate the propagation of firm shocks into worker earnings even in this sample of firms strongly suggests an important role for relationship-specific investments in the generation of labor market rents.

2 Interpreting Wage Fluctuations

In this section, we sketch a simple model of wage determination designed to interpret the propagation of firm-specific productivity shocks into wages. Our model is tailored to the newly-innovative firms that are the focus of our empirical analysis. For the purposes of motivating our model, two features of these firms are notable. First, these firms are relatively small: the median firm in our estimation sample employs 17 workers in the year of its first patent application.¹ Such firms seem unlikely to possess significant market power over new hires or to have reputations that allow them to credibly commit to backloaded compensation schemes. Second, the innovative work conducted at these firms is necessarily specialized and proprietary in nature, likely making it costly to replace incumbent employees with new hires. As in Becker (1964), Stevens (1994), and Manning (2006), the imperfect substitutability of incumbent workers with new hires provides an avenue for incumbent workers to extract rents from the firm in the form of wage premia.

Our model yields a linear wage-setting rule similar to those found in many search models with multi-lateral bargaining (Pissarides 2000; Cahuc and Wasmer 2001; Acemoglu and Hawkins 2014), as well as in much of the classic literature on union wage bargaining (Brown and Ashenfelter 1986). We use this framework to motivate standard empirical “rent-sharing” specifications and to clarify the endogeneity problems that arise when estimating the transmission of firm-specific shocks to wages. We then discuss the assumptions under which patent allowance decisions can facilitate the identification of economic parameters of interest.

2.1 Preliminaries

We work with a one period model. Each firm $j \in \{1, \dots, J\}$ begins the period with I_j incumbent workers and a non-wage amenity value A_j capturing factors such as geographic location and work environment. The firm can hire as many new workers as desired at competitive market wage $w_j^m = w^m(A_j)$. As in classic hedonic models (e.g., Rosen, 1986), the wage demanded by new hires will tend to be decreasing in the value of these amenities (i.e., $\frac{\partial}{\partial A_j} w^m(A_j) \leq 0$), which leads entry wages to vary by firm despite the perfectly competitive nature of this market.

Hiring N_j new workers requires paying a training and recruiting cost $c(N_j, I_j)$. The function $c(\cdot, \cdot)$ exhibits constant returns to scale, which implies

$$c(N_j, I_j) = c(N_j/I_j) I_j.$$

We assume $c(\cdot)$ is twice differentiable and convex.

¹While the firms in our sample are small relative to, e.g., firms included in the Compustat data, they should not be thought of as anomalously small in size. Axtell (2001) finds using Economic Census data that the modal US firm size is 1 employee, and the median is 3 (4 if size 0 firms are not counted).

The firm chooses a wage $w_j^I \geq w_j^m$ for incumbent workers at the beginning of the period. After the wage is posted, incumbent workers receive outside job offers. These offers are non-verifiable, in part because they may involve non-wage amenities, and therefore cannot be matched. However, the firm knows the offers have wage-equivalent values drawn from the following translated Beta distribution

$$G_j(\omega) = \left(\frac{\omega - w_j^m}{\bar{w}_j - w_j^m} \right)^\eta \quad \text{for } \omega \in [w_j^m, \bar{w}_j],$$

where $\bar{w}_j > w_j^m$ is the maximum value of an outside offer. As η grows, offers become concentrated around \bar{w}_j , while, as η tends towards zero, offers become concentrated around w_j^m .

Workers receiving outside offers with value greater than the incumbent wage separate from the firm. Consequently, $G_j(w_j^I) I_j$ incumbent workers are retained for production activities. Note that η can therefore be interpreted as the elasticity of worker retention with respect to the incumbent wage premium $w_j^I - w_j^m$.

At the end of the period, the firm produces $Q_j = T_j L_j$ units of output where $L_j = N_j + G_j(w_j^I) I_j$ gives the number of retained workers and T_j gives the firm's "physical" productivity. Output is sold on a monopolistically competitive product market with inverse product demand $P_j(Q_j) = P_j^0 Q_j^{-1/\varepsilon}$ where $P_j^0 > 0$ is a firm-specific constant capturing the firm's product market power and $\varepsilon > 1$ gives the elasticity of demand. After selling its output and paying the retained workers, the firm shuts down.

2.2 The firm's problem

The firm chooses the number of new hires N_j and an incumbent wage w_j^I to maximize profits. Formally, its problem is to:

$$\max_{\{w_j^I, N_j\}} \underbrace{P_j^0 [T_j (G_j(w_j^I) I_j + N_j)]^{1-1/\varepsilon}}_{\text{revenue}} - \underbrace{c(N_j/I_j) I_j}_{\text{training cost}} - \underbrace{w_j^m N_j - w_j^I G(w_j^I) I_j}_{\text{wage bill}}.$$

At an optimum, the firm equates the marginal cost of a new hire to her marginal revenue product (MRP_j):

$$w_j^m + c'(N_j/I_j) = (1 - 1/\varepsilon) \frac{P(Q_j) Q_j}{L_j} \equiv MRP_j. \quad (1)$$

Note that the marginal cost of a hire exceeds the market wage by the amount of the training/recruiting cost $c'(N_j/I_j)$, which is increasing in the gross hiring rate N_j/I_j .

For incumbent wages, the first order condition can be written:

$$MRP_j = w_j^I + \underbrace{(w_j^I - w_j^m) / \eta}_{\text{inframarginal wage cost}}. \quad (2)$$

As in monopsony models, the firm equates the marginal revenue product of an incumbent worker to her marginal factor cost, which consists of her wage w_j^I plus a term capturing the costs of raising wages for inframarginal incumbents. As the retention elasticity η approaches infinity, this term collapses to the standard neoclassical requirement that the marginal revenue product of an incumbent worker equal her wage.

2.3 Rent Sharing

Subtracting equations (1) and (2), we arrive the following expression for the incumbent wage premium:

$$w_j^I - w_j^m = \frac{\eta}{1 + \eta} c'(N_j/I_j). \quad (3)$$

Incumbents are paid a premium over new hires in proportion to their marginal training/recruiting costs $c'(N_j/I_j)$. When $c'(N_j/I_j) = 0$, incumbent workers are *replaceable*. In this case, the firm views new hires and incumbents as perfect substitutes and pays them equivalently. The fraction $\frac{\eta}{1 + \eta} \in [0, 1]$ plays the role of the exploitation index in classic monopsony models (Manning, 2011) where η would correspond to a firm-specific labor supply elasticity. As the retention elasticity η approaches infinity, incumbents capture their full (marginal) replacement cost in the form of elevated wages. As η tends towards zero, the outside options of incumbents deteriorate, allowing the firm to retain them at the market wage w_j^m and capture the rents in the employment relationship.

Plugging (3) into (1) yields an expression for the incumbent wage that is useful for motivating our empirical rent-sharing specifications:

$$\begin{aligned} w_j^I &= \frac{1}{1 + \eta} w_j^m + \frac{\eta}{1 + \eta} MRP_j \\ &= (1 - \theta) w_j^m + \theta MRP_j \end{aligned} \quad (4)$$

where $\theta = \frac{\eta}{1 + \eta}$. Workers are paid a θ -weighted average of their marginal productivity and the market wage w_j^m . Rewriting $\theta = \frac{w_j^I - w_j^m}{MRP_j - w_j^m}$ illustrates the link to models with Nash wage bargaining in which θ gives the fraction of marginal match surplus paid out in wage premia.² As the retention elasticity η increases, θ rises and workers capture more of the surplus.

The parameter θ has a clear causal interpretation: a dollar increase in marginal productivity yields a θ -cent pay increase for incumbents. It is useful to review briefly why marginal products can vary in this model. In the special case where incumbents are replaceable ($c'(N_j/I_j) = 0$), (1) implies the marginal revenue product would be pinned

²Stole and Zwiebel (1996) propose a multilateral bargaining framework where workers and firms also bargain over infra-marginal products. This bargaining concept is embedded in a search and matching framework by Acemoglu and Hawkins (2014). Given our assumption of a constant product demand elasticity, the wage rule that results from the Stole-Zwiebel approach is analogous to equation (4) with the modification that the weights on the reservation wage and marginal revenue product need not sum to 1.

to the market wage w_j^m . Hence, there would be no scope for fluctuations in MRP_j other than due to shifts in the amenity vector A_j . But when incumbents are not replaceable, MRP_j will also respond to fluctuations in “revenue productivity” $P_j^0 T_j$. As described below, our empirical approach uses variation in patent allowances to isolate the variation in MRP_j that arises due to exogenous fluctuations in revenue productivity.

2.4 Estimating pass-through

We can operationalize equation (4) by plugging in the definition of MRP_j to get:

$$\begin{aligned} w_j^I &= (1 - \theta) w_j^m + \theta \left(1 - \frac{1}{\varepsilon}\right) \frac{P_j Q_j}{L_j} \\ &= (1 - \theta) w_j^m + \pi S_j. \end{aligned} \tag{5}$$

The last line of this expression is a standard empirical rent-sharing specification relating incumbent wages at the firm to a measure of average labor productivity $S_j = \frac{P_j Q_j}{L_j}$, which we refer to as *gross surplus* per worker.

The parameter $\pi = \theta \left(1 - \frac{1}{\varepsilon}\right)$ governs pass-through of gross surplus to wages and can be thought of as the labor market analog of cost-price pass-through coefficients often used to study imperfect competition in product markets (Goldberg and Hellerstein 2013; Weyl and Fabinger 2013; Gorodnichenko and Talavera 2017). The term $\left(1 - \frac{1}{\varepsilon}\right)$ is an adjustment factor that converts average labor productivity to marginal labor productivity. While π is our primary parameter of interest, we also explore calibrations of ε and consider the implied values of the structural rent-sharing coefficient θ .

Card et al. (2018) review several studies that use panel methods to assess the relation between the wage growth of incumbent workers and fluctuations in various measures of firm surplus. Equation (5) suggests such specifications will suffer from omitted variables bias whenever surplus fluctuations are correlated with changes in the market wage w_j^m . For example, shocks to firm productivity may contain a market-wide component. If all firms in a market become more productive, market wages will rise. This possibility would lead to a misattribution of market-level wage adjustments to rent-sharing and a corresponding upward bias in OLS estimates of π .

A different class of potential biases arises from unobserved shocks to the amenity value of a firm. Suppose the work environment at a firm improves and leads to a decrease in w_j^m . This improvement will lead, *ceteris paribus*, to an increase in firm scale, which will tend to depress average labor productivity through drops in the product price $P_j(Q_j)$. Consequently, such shocks will induce a positive covariance between w_j^m and $S_j = P_0 T_j^{1-1/\varepsilon} L_j^{-1/\varepsilon}$ and hence lead to an overstatement of the degree of rent-sharing. However, unobserved amenity shocks could also exert a direct effect on productivity. For example, a recent empirical literature finds that variation in management practices affects both worker morale and productivity (Bloom and Van Reenen 2007; Bender et al. 2018). A new manager who motivates workers could plausibly raise total factor productivity T_j while lowering the market

wage w_j^m via increases in the amenity value A_j of the firm. This possibility would lead to an under-estimate of rent-sharing as the productivity shock is accompanied by an unobserved amenity shock.

2.5 Instrumenting with Patent Decisions

To circumvent these endogeneity problems, we use the initial decision of the US Patent and Trademark Office (USPTO) on a firm's first patent application as an instrument for the firm's surplus.³ Patents could influence average labor productivity through two channels, both of which provide valid identifying variation. First, a patent grant could raise a firm's product price intercept P_j^0 by creating a barrier to competition by rival firms.⁴ Second, a patent grant could raise a firm's TFP T_j by making it profitable for the firm to implement the patented technology.

We document below that within observable strata, the USPTO's initial decision on a given patent application is unrelated to trends in firm performance, implying that initial patent decisions are as good as randomly assigned with respect to counterfactual *changes* in firm outcomes. Consistent with this evidence, we also document below that it is hard to predict initial decisions using firm characteristics in the year of application. Finally, we assume that patent decisions are uncorrelated with fluctuations in the market wage w_j^m . In the model above, this condition is sufficient to ensure that instrumenting S_j with the patent allowance isolates exogenous variation in revenue productivity $P_j^0 T_j$ and identifies the pass-through parameter π .

The assumption that patent decisions are uncorrelated with fluctuations in w_j^m merits further discussion in our setting as several violations of this condition are conceivable, most of which are not explicitly modeled in the above framework. A first concern is that patent allowances might lead the firm to demand more hours from workers, in which case w_j^m would rise. However, we would expect this to be a short-run phenomenon that dissipates as the firm expands towards its new target size, and we find no evidence of such wage dynamics in the data. A different sort of violation would occur if patents shift expectations about firm growth and therefore about the future earnings growth of workers. This sort of mechanism arises in dynamic wage posting models with offer matching (Postel-Vinay and Robin 2002) and would imply that w_j^m falls in response to an allowance. However, such a violation would also imply that initial allowances should raise the wage growth of new hires, an assertion for which we find no empirical support.

A second concern is that initial allowance decisions might be geographically correlated, in which case instrumenting with initial allowances might pick up market-wide fluctuations in w_j^m . We show below however that

³Van Reenen (1996) also investigated patents as a source of variation, but found them to be a relatively weak predictor of firm profits in his sample of firms (see his footnote 11). This finding is in keeping with the notion that most patents generate little ex-post value to the firm (Pakes 1986), motivating our focus on ex-ante valuable patent applications as described in Section 5. A natural alternative empirical strategy in our setting would be to use the leniency of the patent examiner assigned to review the patent application as an instrument, as in Sampat and Williams (forthcoming). Unfortunately, this strategy reduces the precision of our estimates to the point of being uninformative.

⁴Perhaps the classic example is patents on branded small molecule pharmaceuticals. In the absence of patents, many branded pharmaceuticals would experience near-immediate entry of generic versions which compete with branded pharmaceuticals at close to marginal cost prices.

the intra-class correlation of initial patent allowances within geography and sector is indistinguishable from zero, which suggests that allowances are best thought of as truly firm-specific shocks. We also find no impact of patent allowances on the earnings of workers in their first year of employment with a firm, which should provide a reasonable proxy of the market wage w_j^m . Since all of the above concerns involve correlations between patent allowances and fluctuations in the market wage w_j^m , this provides a strong corroboration of the exogeneity of the patent allowance instrument.

A final concern is that firms may respond to patent decisions by changing the composition of their workforce. By leveraging the panel structure of our data, we can directly investigate whether firms change their composition of new hires (or separations) in response to patent allowances. We also address this concern by analyzing the wage *growth* of incumbent workers, which by construction differences out any selection on time invariant characteristics. In practice, we find that such adjustments have little effect on our estimates of the pass-through parameter π .

3 Data and Descriptive Statistics

To conduct our empirical analysis, we construct a novel linkage of several administrative databases, which provides us with panel data on the patent filings, patent allowance decisions, and outcomes of US firms and workers.

3.1 USPTO Patent Applications

We begin with public-use administrative data on the universe of patent applications submitted to the US Patent and Trademark Office (USPTO) since late 2000.⁵ We link these published US patent applications with several USPTO administrative datasets. Because published patent applications are not required to list the assignee (owner) of the patent, approximately 50% of published patent applications were originally missing assignee names. We worked with the USPTO to gain access to a separate public-use administrative data file that allows us to fill in assignee names for most of these applications. The public-use USPTO PAIR (Patent Application Information Retrieval) administrative data records the full correspondence between the applicant and the USPTO, allowing us to infer the timing and content of the USPTO’s initial decision on each patent application as well as other measures of USPTO and applicant behavior. Details on these and the other patent-related data files that we use are included in Appendix A.

Panel A of Table 1 describes the construction of our patent application sample. Our full sample consists of

⁵The start date of our sample is determined by the American Inventors Protection Act of 1999, which required publication of nearly the full set of US patent applications filed on or after 29 November 2000. We say “nearly” because our sample misses patent applications that opt out of publication; Graham and Hegde (2014) use internal USPTO records to estimate that around eight percent of USPTO applications opt out of publication.

the roughly 3.6 million USPTO patent applications filed on or after 29 November 2000 that were published by 31 December 2013; we restrict attention to applications filed on or before 31 December 2010 in order to limit the impact of censoring. We drop around 400,000 applications that are missing assignee names and therefore cannot be matched to business tax records. We also limit our sample to standard (so-called “utility”) patents.⁶

To focus on a subset of firms for which patent allowances are most likely to induce a meaningful shift in firm outcomes, we make several restrictions that aim to limit our sample to first-time patent applicants. First, we drop so-called “child” applications that are derived from previous patent applications. Second, we retain the earliest published patent application observed for each assignee in our sample.⁷ Finally, we exclude assignees which we observe to have had patent grants prior to the start of our published patent application sample.⁸ Ideally, we would exclude assignees that had patent *applications* (not just patent grants) prior to the start of our published patent application sample, but unsuccessful patent applications filed before 29 November 2000 are not publicly available. These restrictions leave a sample of around 96,000 patent applications, which we then attempt to match to our US Treasury business tax files.

3.2 Treasury Tax Files

We link US Treasury business tax filings with worker-level filings. Annual business tax returns record firm outcomes from Form 1120 (C-Corporations), 1120S (S-Corporations), and 1065 (Partnership) forms, and cover the years 1997-2014. The key variables that we draw from the business tax return filings are revenue, value added, EBITD (earnings before interest, taxes, and deductions), and labor compensation; each of these is defined in more detail in Appendix A.

We link these business tax returns to worker-level W2 and 1099 filings in order to measure employment and compensation for employees (e.g., wage bill) and independent contractors, respectively, at the firm-year level. The relevant variables are defined in more detail in Appendix A. We winsorize all monetary values in the tax files from above and below at the five percent level, which is standard when working with the population of US Treasury business tax files (see, for example, DeBacker et al. 2016; Yagan 2015). Since our analysis focuses on per-worker outcomes, we winsorize outcomes on a per-worker basis.

To distinguish employment and compensation for inventors and non-inventors, we use Bell et al.’s (2016) merge of inventors listed in patent applications to W2 filings. Inventors are defined as individuals ever appearing

⁶Utility patents, also known as “patents for invention,” comprise approximately 90% of USPTO-issued patent documents in recent years; see <https://www.uspto.gov/web/offices/ac/ido/oeip/taf/patdesc.htm> for details.

⁷Because USPTO procedure assigns application numbers sequentially, we break ties in the cases in which a given assignee submits multiple applications on the same day by taking the smallest application number.

⁸We search for such patent grants going back to 1976, the date with electronic patent grant records are most easily available. Given the firms in our sample, the likelihood that a firm had a patent granted prior to 1976 seemed sufficiently small not to warrant a more extensive attempt to match to earlier patent grants.

in the Bell et al. (2016) patent application-W2 linkage, rather than individuals listed as inventors on the specific patent application relevant to a given firm in our sample.

3.3 Linkage Procedure

We build on the name standardization routine used by the National Bureau of Economic Research (NBER)'s Patent Data Project (<https://sites.google.com/site/patentdataprotect/>) to implement a novel firm name-based merge of patent assignees to firm names in the US Treasury business tax files. Specifically, we standardize the firm names in both the patent data and (separately) the US Treasury business tax files in order to infer that, e.g., "ALCATEL-LUCENT U.S.A., INC.," "ALCATEL-LUCENT USA, INCORPORATED," and "ALCATEL-LUCENT USA INC" are all in fact the same firm. We then conduct a fuzzy merge of standardized assignee names to standardized firm names in the business tax files using the SoftTFIDF algorithm based on a Jaro-Winkler distance measure. This merge is described in more detail in Appendix A.

To assess the quality of our merge, we conduct two quality checks: first, we validate against a hand-coded sample; and second, we validate against the inventor-based linkage of Bell et al. (2016). As described in Appendix A, the results of these validation exercises suggest that our merge is of relatively high quality, with type I and II error rates on the order of five percent.

Panel B of Table 1 describes our linkage between the USPTO patent applications data and the US Treasury business tax files. Of the around 96,000 patent applications we attempt to match to the US Treasury business tax files, we match around 40,000 patent applications. The USPTO estimates that in 2015 approximately 49.6% of USPTO patent grants were filed by US-based assignees, which implies our match rate to US-tax-paying entities is on the order of 83%.⁹ These 40,000 patent applications are matched to around 40,000 standardized firm names in the US Treasury business tax files, which correspond to 82,000 firms (employer identification numbers, or EINs).

We build the analysis sample from these 82,000 EINs in four steps. Our goal here is to construct a unique and well-defined match between patent applications and firms in a subset of firms for which patent allowances are most likely to induce a meaningful shift in firm outcomes. First, we attempt to restrict our post-merge tax analysis sample to first-time patent applicants by retaining the earliest-published patent application observed for each EIN, and by excluding EINs which we observe to have had patent grants prior to the start of our published patent application sample. Second, in cases in which there are multiple EINs for a standardized name in the tax files, we keep the EIN with largest revenue in the year that the patent application was filed. Third, we restrict attention to "active" firms, defined as EINs that have a positive number of employees in the year of application and non-zero, non-missing total income or total deductions in the year the patent application was filed and in the three previous years. This

⁹These USPTO estimates, which are based on the reported location of patent assignees, are available here: https://www.uspto.gov/web/offices/ac/ido/oeip/taf/own_cst_utl.htm.

restriction allows us to investigate pre-trends in our outcome variables among economically relevant firms. Fourth, we limit attention to EINs with less than 100 million in revenue in 2014 USD in the year of patent application. This step, which eliminates firms in the top centile of the firm size distribution, allows us to avoid complexities related to the largest multinational companies and focus on firms for whom patent allowance decisions are more likely to be consequential.¹⁰ These restrictions leave us with a sample of 9,732 patent applications, each uniquely matched to one EIN in the US Treasury business tax files. It is worth noting that focusing on such a small subset of firms is common in analyses such as ours. For example, Kogan et al. (2017) start with data on 7.8 million granted patents, which they winnow down to a final sample of 5,801 firms with at least one patent.

3.4 Measuring Surplus

As described in Card et al. (2018), empirical rent-sharing estimates are often sensitive to a number of measurement issues, the most prominent of which is the choice of rent measure. In keeping with equation (5), we rely on a *gross surplus* measure of rent that differs from “match surplus” due to the absence of data on workers’ reservation wages. Letting Π_j denote the firm’s economic profits, the model of Section 2 implies the firm’s total gross surplus can be written:

$$S_j L_j = \underbrace{w_j^m N_j + w_j^l G(w_j^l) I_j}_{\text{wage bill}} + \underbrace{\Pi_j}_{\text{profits}} + \underbrace{c(N_j/I_j) I_j}_{\text{training/recruiting costs}} .$$

To measure this theoretical concept in the tax data, we use the sum of the firm’s W2 earnings in a year and its earnings before interest, tax, and depreciation (EBITD). Though firms sometimes report negative EBITD, this surplus measure is usually positive and provides a plausible upper bound on the flow of resources capable of being captured by workers. Note that this measure is theoretically justified by the presumption that firms do *not* claim deductions on training costs; i.e., that EBITD captures the sum $\Pi_j + c(N_j/I_j) I_j$.¹¹

For comparison with past work, we also report results that use a value added measure of surplus. Our approximation to value added comes from line 3 of Form 1120, which deducts from gross sales, returns and allowances, and the cost of goods sold. This measure suffers from the disadvantage that it may include a number of additional unobserved firm costs including rents, advertising, and financing fees that are likely unavailable for capture by workers.

¹⁰Statistics for firm size distribution are from Smith et al. (2017). Specifically, in the full population of C-corporations, S-corporations, and Partnerships with positive sales and positive W2 wage bills, \$100 million in revenue in 2014 USD falls in the top one percent of firms.

¹¹Unlike some other capital expenses and costs related to intangibles, which can be amortized, firms typically can not amortize and deduct costs related to training. Specifically, section 197 on intangibles includes workforce in place (e.g., “experience, education, or training”) and business books and records (e.g., “intangible value of technical manuals, training manuals or programs”) in the list of assets that cannot be amortized for most firms. See <https://www.irs.gov/pub/irs-pdf/p535.pdf> for additional details.

3.5 Summary Statistics

Table 2 tabulates summary statistics on our firm and worker outcomes in each of two samples: our analysis sample of matched patent applications/firms (N=9,732), and our sub-sample of matched patent applications/firms for which the patent applications are in the top quintile of predicted value (N=1,946), which will be defined in the next section. All summary statistics are as of the year the patent application was filed.

Panel A documents summary statistics on firm-level outcomes. In our analysis sample, the median firm generated around three million dollars in revenue, employed 17 workers, and reported roughly \$7,000 in EBITD per worker. Approximately 8% of patent applications were initially allowed. Panel B documents summary statistics on worker-level outcomes. The median firm in our analysis sample paid \$48,000 in annual earnings per W2 employee, employed a workforce that was approximately 75% male, and issued 2.5% of its W2s to individuals listed as inventors on at least one patent application. Contract work turns out to be relatively uncommon in this sample, with 1099s constituting only about 10% of the sum of W2 and 1099 employment for the median firm.

4 Institutional Context: Initial Patent Decisions

The US Patent and Trademark Office (USPTO) is responsible for determining which — if any — inventions claimed in patent applications should be granted a patent. Patentable inventions must be patent-eligible (35 U.S.C. §101), novel (35 U.S.C. §102), non-obvious (35 U.S.C. §103), useful (35 U.S.C. §101), and the text of the application must satisfy the disclosure requirement (35 U.S.C. §112). When patent applications are submitted to the USPTO, they are routed to a central office which directs the application to an appropriate “art unit” that specializes in the technological area of that application. For example, art unit 1671 reviews applications related to the chemistry of carbon compounds, whereas art unit 3744 reviews applications related to refrigeration. The manager of the relevant art unit then assigns the application to a patent examiner for review. If the examiner issues an initial allowance, the inventor can be granted a patent. If the examiner issues an initial rejection, the applicant has the opportunity to “revise and resubmit” the application, and the applicant and examiner may engage in many subsequent rounds of revision (see Williams 2017 for more details).

Our empirical strategy focuses on contrasting firms that receive an initial allowance to other firms that applied for a patent but received an initial rejection. Empirically, most patent applications receive an initial decision within three years of being filed (see Appendix Figure C.1). While some applications that are initially rejected receive a patent grant relatively quickly, the modal application that is initially rejected is never granted a patent (see Appendix Figure C.2).

Because our empirical strategy will contrast firms whose applications are initially allowed to those whose applications are initially rejected, having some sense of what predicts initial allowance decisions is useful. Table

3 reports least squares estimates of the probability of an initial allowance as a function of firm characteristics in the year of application. Column (1a) shows that predicting initial allowances is surprisingly difficult. Applications from firms with more W2 employees are somewhat less likely to be initially allowed, as are those from firms with higher value added per worker. Jointly, the covariates are statistically significant. Column (1b) adds art unit by application year fixed effects that control for technology-specific changes over time. This simple addition renders all baseline covariates statistically insignificant both individually and jointly, which provides some assurance that initial patent decisions are not strongly dependent on baseline firm performance. Given this empirical evidence, we proceed by assuming that any remaining selection is on time-invariant firm characteristics that can be captured by firm fixed effects.

A separate concern has to do with whether initial allowances are best thought of as idiosyncratic or market-level shocks. Seminal work by Jaffe, Trajtenberg, and Henderson (1993) demonstrated that patent citations are highly localized geographically. To test whether initial allowances are also geographically clustered, we fit linear random effects models to the initial allowance decision. Appendix Table C.1 reports intraclass correlations at various levels of geography before and after subtracting off art unit by application year mean allowance rates. In either case, the within-state correlation is estimated to be zero, while the correlation within five-digit ZIP codes is quite low (0.07-0.10) and statistically indistinguishable from zero. These findings indicate that initial allowances are best thought of as truly idiosyncratic firm-specific shocks that are unlikely to elicit market-wide wage responses.

5 Detecting Valuable Patents

The value distribution of granted patents is heavily skewed (Pakes 1986), which suggests that low-value patent applications — if granted — are unlikely to generate meaningful shifts in firm outcomes. Constructing a measure of the ex-ante value of patent applications enables us to focus our analysis on patent applications that are likely to induce changes in firm behavior.

A variety of metrics have been proposed as measures of the value of granted patents, including forward patent citations (Trajtenberg 1990), patent renewal behavior (Pakes 1986; Schankerman and Pakes 1986; Bessen 2008), patent ownership reassignments (Serrano 2010), patent litigation (Harhoff, Scherer, and Vopel 2003), and excess stock market returns (Kogan et al. 2017). These value measures encounter three challenges in our empirical context. First, these measures are only defined for granted patents, whereas we would like to take advantage of data on patent applications, including those that are ultimately unsuccessful. Second, most of these measures arguably correspond to a measure of social value — or social spillovers, in the sense of social value minus private value — whereas we are more interested in measuring firms' private value of a patent. This issue arises most sharply with forward patent citations, which are typically used as a measure of spillovers (e.g., Bloom, Schankerman, and Van Reenen

2013). Third, all of these measures are defined ex-post: citations, renewals, reassignments, and litigation are often measured many years after the initial patent award. But in our context — as in Kogan et al. (2017) — what is arguably more relevant is the expected private value of the patent at the time of the patent application or patent grant.

To this end, we build on the recent analysis of Kogan et al. (2017) (henceforth, KPSS), who measure the high-frequency response of stock prices around the date of patent grant announcements to estimate the value of patent grants that are awarded to publicly traded companies. We estimate a simple statistical model designed to extrapolate their estimates to non-publicly traded companies and to non-granted patent applications in our analysis sample.

We model the KPSS patent value ξ_j for each firm-patent application j in our data as obeying the following conditional mean restriction:

$$E[\xi_j | X_j, \mathbb{A}_j] = \exp(X_j' \delta + v_{\mathbb{A}_j}),$$

where X_j denotes a vector of baseline firm and patent application covariates and \mathbb{A}_j denotes the art unit to which the application was assigned. The exponential functional form underlying this specification is designed to accommodate the fact that the KPSS values are non-negative and heavily skewed. Because we have, on average, only 2.3 applications with non-missing ξ_j per art unit, some penalization is required to avoid overfitting. Accordingly, we treat the art unit effects $\{v_a\}$ as *i.i.d.* draws from a normal distribution with unknown variance σ_v^2 rather than fixed parameters to be estimated. The model is fit via a random effects Poisson maximum likelihood procedure. As described in Appendix B, this procedure exploits the conditional mean restriction $E[\xi_a | X_a] = \int \exp(X_a' \delta + v) \omega_a(v) dv$ where ξ_a is the vector of KPSS values in an art unit a , X_a is the corresponding vector of baseline application and firm predictors, and $\omega_a(v)$ is the posterior distribution of v_a given the observed data (ξ_a, X_a) .

Table 4 reports the Poisson parameter estimates. Applications submitted to more countries (“patent family size”) tend to be of higher value, as do applications with more claims and applications submitted by firms with larger revenues.¹² We also document substantial variability of patent value across art units: a standard deviation increase in the art unit random effect is estimated to raise mean patent values by 127 log points. This variability finding is of interest in its own right as it suggests that patent decisions involve much higher stakes in some USPTO art units than others.

We use our estimates of the parameters (δ, σ_v) to compute Empirical Bayes predictions $\hat{\xi}_j$ of ξ_j for every

¹²The number of countries to which an application was submitted, often referred to as patent “family size,” is defined as a set of patent applications filed with different patenting authorities (e.g., the US, Europe, Japan) that refer to the same invention; work starting with Putnam (1996) has argued that firms should be willing to file more privately valuable patents in a larger number of countries. Patents list “claims” over specific pieces of intellectual property, and work starting with Lanjouw and Schankerman (2001) has argued that patents with a larger number of claims may be more privately valuable. See Appendix A for details on both these measures.

patent application in our analysis sample, including those that lack a KPSS value either because the application is assigned to a privately held firm, or because the application is never granted a patent.¹³ Empirically, these predictions are highly accurate: a least squares fit of ξ_j to $\hat{\xi}_j$ yields a slope of 1.12 and an R^2 of 68%. Figure 1 shows that binned average KPSS values track the Empirical Bayes predictions very closely. Appendix Table C.2 lists mean predicted values by subject matter area.

The ultimate test of $\hat{\xi}_j$ is whether it predicts treatment effect heterogeneity: that is, do allowances of patent applications of higher predicted value result in larger shifts in firm outcomes? To investigate this question, we fit a series of interacted difference-in-differences models of the following form:

$$Y_{jt} = \alpha_j + \kappa_{t,k(j)} + Post_{jt} \cdot \left[\sum_{b=1}^5 s_b(\hat{\xi}_j) \cdot (\tilde{\psi}_b + \tilde{\tau}_b \cdot IA_j) \right] + r_{jt} \quad (6)$$

where Y_{jt} is an outcome for firm (EIN) j in year t , α_j are firm fixed effects, and $\kappa_{t,k(j)}$ are calendar year fixed effects that vary by art-unit/application year cell $k(j)$. The variable $Post_{jt}$ is an indicator for having received an initial patent decision, IA_j is an indicator for whether the patent application is initially allowed, and $\{s_b(\cdot)\}_{b=1}^5$ is a set of basis functions defining a natural cubic spline with five knots.¹⁴ Intuitively, this specification compares initially allowed and initially rejected applications in the same art unit by application year cell, before and after the date of the initial decision. The spline interactions allow the effects of an initial allowance to vary flexibly with the predicted patent value $\hat{\xi}_j$.

Of primary interest is the “dose-response” function $d(x; \tilde{\tau}) \equiv \sum_{b=1}^5 s_b(x) \tilde{\tau}_b$, which gives the effect of an initial allowance for a patent with predicted value x . Figure 2 plots our estimates of this function for a grid of values x when Y_{jt} is either surplus per worker or wage bill per worker. In both cases, we find evidence of an S-shaped response: impacts of initial allowances on both wages and surplus are small and statistically insignificant at low predicted value levels, corroborating both the exclusion and random assignment assumptions underlying our research design. Patents with ex-ante predicted patent values above \$5 million in 1982 USD — roughly the 80th percentile of the predicted value distribution — have larger, statistically significant treatment effects that increase rapidly before stabilizing at values near \$12 million in 1982 USD.¹⁵

Given the S-shaped pattern of treatment effect heterogeneity documented in Figure 2, our empirical analysis pools the bottom four quintiles together and focuses on estimating the impacts of patents in the top quintile of ex-ante predicted patent value. Reassuringly, Columns (2a) and (2b) of Table 3 show that initial allowances are

¹³In cases where no valid KPSS values are present in the entire art unit, we form our prediction by imputing an art unit random effect of zero.

¹⁴The natural cubic spline is a cubic b-spline that imposes continuous second derivatives everywhere but allows the third derivative to jump at the knots (see Hastie, Tibshirani, and Friedman 2016 for discussion). Following Harrell (2001), we space knots equally at the 5th, 27.5th, 50th, 72.5th, and 95th percentiles of the distribution of patent values, which correspond to dollar values of roughly \$0.1M, \$0.7M, \$1.7M, \$4.1M, and \$19.0M 1982 USD respectively. The spline is constrained to be linear below the 5th and above the 95th percentiles.

¹⁵We reference 1982 dollars because those are the units used by KPSS.

equally difficult to predict with baseline characteristics within the top quintile of predicted value, especially after art unit by application year fixed effects have been included. Likewise, Columns (3a)-(4b) of Appendix Table C.1 show that among top-quintile applications, initial allowances continue not to exhibit spatial correlation.

6 Reduced Form Estimates

The treatment effect heterogeneity documented in Figure 2 demonstrates that firms experience economically and statistically significant increases in profitability and wages when valuable patent applications are allowed. However, a natural concern is that these findings could reflect pre-existing trends rather than causal effects of the patent decisions themselves. To investigate this concern, we estimate a series of “event study” specifications of the following form:

$$Y_{jt} = \alpha_j + \kappa_{t,k(j)} + Q5_j \cdot \left[\sum_{m \in \mathcal{M}} D_{jt}^m \cdot (\psi_{5,m} + \tau_{5,m} \cdot IA_j) \right] + (1 - Q5_j) \cdot \left[\sum_{m \in \mathcal{M}} D_{jt}^m \cdot (\psi_{<5,m} + \tau_{<5,m} \cdot IA_j) \right] + r_{jt} \quad (7)$$

where $Q5_j$ is an indicator for the firm’s patent application being in the top quintile of predicted ex-ante value, D_{jt}^m is an indicator for firm j ’s decision having occurred m years ago, and the set $\mathcal{M} = \{-5, -4, -3, -2, 0, 1, 2, 3, 4, 5\}$ defines the five-year horizon over which we study dynamics.¹⁶ The coefficients $\{\hat{\psi}_{5,m}, \hat{\psi}_{<5,m}\}_{m \in \mathcal{M}}$ summarize trends in mean outcomes relative to the date of an initial decision, which may differ by the firm’s ex-ante patent value quintile. Of primary interest are the coefficients $\{\hat{\tau}_{5,m}, \hat{\tau}_{<5,m}\}_{m \in \mathcal{M}}$, which summarize the *differential* trajectory of mean outcomes for initially allowed and initially rejected firms by time relative to the initial decision for top-quintile and lower-quintile value observations, respectively.

Figure 3 plots the coefficients $\{\hat{\tau}_{5,m}, \hat{\tau}_{<5,m}\}_{m \in \mathcal{M}}$ from equation (7) for our main firm outcome variable, surplus. The estimated coefficients illustrate that, among firms with patent applications in the top quintile of the predicted value distribution, firms whose applications are initially allowed exhibit similar trends in surplus per worker to those whose applications are initially rejected in the years prior to the initial decision. However, surplus per worker rises differentially for allowed firms in the wake of an initial allowance, and remains elevated afterwards. Firms with lower predicted value applications, by contrast, exhibit no detectable response of surplus per worker to an initial allowance. Figure 4 documents similar patterns in our main worker outcome variable, wage bill per worker. As expected, the wage response to an initial allowance is muted relative to the surplus response; the ratio of these two impacts provides a crude estimate of the pass-through coefficient π of roughly one-third.

¹⁶We “bin” the endpoint dummies so that D_{jt}^5 is an indicator for the decision having occurred five or more years ago and D_{jt}^{-5} is an indicator for the decision being five or more years in the future.

While wages and surplus respond rather immediately to top-quintile initial allowances, Figure 5 reveals that firm size (as measured by the log number of employees) responds more slowly, taking roughly three years to scale to its new level. The fact that earnings impacts remain stable over this horizon casts doubt on the possibility that the impacts in Figure 3 are driven primarily by an increase in hours worked (which we cannot observe in tax data) rather than an increase in hourly wages. The nearly immediate response of surplus and wages to initial allowances may signal that our panel of relatively small innovative firms was initially credit constrained. Evidence from Farre-Mensa, Hegde, and Ljungqvist (2017), who document that patent grants are strongly predictive of access to venture capital financing, corroborates this view. Access to venture capital and other forms of financing is a plausible additional channel through which patent decisions could quickly affect the marginal revenue product of labor and consequently worker wages.¹⁷

As background for interpreting the magnitude of these results, Figure 6 documents that an initial allowance raises the probability of having the patent application granted by roughly 50% in the year after the decision, with gradual declines afterwards. The probability of receiving a patent grant jumps by less than 100% for two reasons. First, some initially allowed applications are not pursued by applicants, possibly because the assignee went out of business while awaiting the initial decision, or because the applicant learned new information since filing which led them to believe that the patent was not commercially valuable. Second, as described in Section 4, many initially denied applications reapply and eventually have their applications allowed. Our estimates in Figure 6 suggest that the impact of initial allowances on patent grants is somewhat smaller for higher-value patents, perhaps because they are more likely to be approved shortly after a rejection; a pooled difference-in-difference estimate of the impact on the grant probability of high-value patents is approximately one third. Hence, the impact of high-value patent grants on firm outcomes is likely to be roughly three times the impact of an initial allowance on firm outcomes, though it is possible that allowances influence firm outcomes independent of grant status if allowances relieve credit constraints before a patent has actually been granted. In what follows, we continue to report the reduced form impacts of allowances as our ultimate goal is to instrument for surplus rather than for patent grants.

¹⁷As a robustness check we fit a version of (7) allowing linear interactions of D_{jt}^0 and $IA_j \cdot D_{jt}^0$ with the week of the patent decision. We find that the contemporaneous surplus impacts we observe are increasing in the number of days that have elapsed since the initial decision. We find no contemporaneous effect of initial patent allowances decided in late December on either surplus or wages, which reassures us that the effect is not abnormally immediate.

6.1 Impacts on Firm Averages

Table 5 pools pre- and post-application years and quantifies the average effects displayed in the event study figures by fitting simplified difference-in-differences models of the following form:

$$Y_{jt} = \alpha_j + \kappa_{t,k(j)} + Q5_j \cdot Post_{jt} \cdot (\psi_5 + \tau_5 \cdot IA_j) + (1 - Q5_j) \cdot Post_{jt} \cdot (\psi_{<5} + \tau_{<5} \cdot IA_j) + r_{jt}. \quad (8)$$

The parameters reported in Table 5 are τ_5 and $\tau_{<5}$, which respectively govern the effects of top-quintile and lower-quintile value patents being initially allowed.

Column 1 of Table 5 documents that initial allowances have no effect on the probability of firm survival, as proxied by the presence of at least one W2 employee. Given this result, the remainder of the columns in this table focus on outcomes conditional on firm survival as measured by the presence of at least one W2 employee (hence the smaller sample sizes in subsequent columns). Column 2 of Table 5 reports the impact of an initial allowance on the log of firm size, as measured by the number of W2 employees at the firm.¹⁸ Having a top-quintile patent allowed leads the firm to expand by roughly 22%. Notably, initial allowances of patents with lower predicted value have no detectable impact on firm survival, firm size, or any other outcome that we examine; these results suggest that differential trends for initially allowed and initially rejected patents are unlikely to confound our analysis.

An allowance of a high-value patent application is associated with roughly \$37,000 in additional revenue per worker (Column 3 of Table 5) and roughly \$16,000 in value added per worker (Column 4 of Table 5). EBITD per worker rises by roughly \$9,000 (Column 5 of Table 5), which we interpret as income to firm owners, while wage bill per employee rises by roughly \$3,600 (Column 6 of Table 5). Our surplus measure, which sums EBITD and wage bill, rises by \$12,400 per worker (Column 7 of Table 5).¹⁹ As described in Section 3, we interpret our estimated effects on surplus as the impact on total operating cash flow at the firm. In this paper, our central interest is in estimating how this surplus measure is divided between workers and firm owners.

Table 5 also reports impacts on various measures of labor compensation. A successful top-quintile patent application is associated with an increase in firm-level deductions for labor-related expenses of around \$4,000 (Column 8 of Table 5), which is roughly comparable to what we found for wage and salary compensation based on W2 wage bills. On the other hand, pooling W2 earnings with 1099 earnings yields an impact of only \$2,800 per worker (Column 9 of Table 5). In percentage terms, these impacts are fairly close: labor compensation per W2 rises by roughly 7.7%, while W2 + 1099 earnings per W2 rise by 6.2%. However, these results suggest that 1099

¹⁸We work with logarithms for firm size because this variable is not winsorized and is very heavily skewed.

¹⁹The sum of the per-worker impacts on EBITD and wage bill does not exactly match the impact on surplus per worker because the variables are winsorized separately.

compensation is, if anything, less responsive to shocks than W2 wages and salaries.

Finally, the last column of Table 5 reports impacts on a measure of the average individual income tax burden per worker.²⁰ An initial allowance of a high-value patent is estimated to yield \$770 of additional tax revenue per worker. Although this figure is statistically indistinguishable from zero, the point estimate implies an effective marginal tax rate of 21% on the \$3,600 of extra W2 earnings reported in Column 6 of Table 5, which is roughly the average US marginal tax rate found in TAXSIM (see Feenberg and Coutts 1993) over our sample period.²¹ In percentage terms, an initial allowance of a high-value patent raises tax revenue per worker by 5% — slightly below the proportional impact on W2 earnings per worker. This finding suggests the presence of an important fiscal externality between corporate tax treatment of innovation and income tax revenue.²²

Panel B of Table 5 repeats the above impact analysis on the subset of “closely held” firms registered as partnerships or S-corporations. Because these businesses rarely offer stock compensation, wage responses are likely to provide a more comprehensive measure of rent sharing in this subsample (see Smith et al. 2017). Among closely-held firms we find somewhat larger impacts on revenue, value added, and EBITD per worker accompanied by commensurately large impacts on average wages and labor compensation. In our pooled sample the ratio of the impact on wage bill per worker to the impact on surplus per worker is 29 cents, whereas the ratio at closely-held firms is 27 cents; the close similarity of these two estimates suggests that the inability to offer stock options does not dramatically alter the pass-through from firm-specific shocks to worker wages. Appendix Table C.3 shows that patent allowances also have similar effects on firms in the top and bottom half of the distribution of initial firm sizes.

6.2 Impacts on Workforce Composition

A difficulty with interpreting impacts on firm-level aggregates is that firms may alter the skill mix of their employees in response to shocks, in which case changes in wages could simply reflect compositional changes rather than changes in the compensation of similar employees. Van Reenen (1996, pp. 216-217) provided a back-of-the-envelope calculation suggesting that compositional changes were unlikely to be a major concern in his sample. In Table 6 we directly investigate the possibility of such compositional changes using our link of W2’s to EINs.

Columns 1 and 2 of Table 6 reveal that neither the share of employees who are women nor the share of

²⁰Our measure, which is the main tax variable in the databank (the main panel dataset used by researchers using the US Treasury tax files), captures “tentative” tax burden before accounting for the Alternative Minimum Tax. It is not available in a small number of cases, which is why Column 10 has slightly fewer observations than the per W2 worker columns.

²¹See <http://users.nber.org/~taxsim/all1up/all1.html> for annual estimates.

²²One specific implication of this finding is that patents influence the revenue raised from both business and individual income taxes. Consequently, so-called “patent box” proposals, which are designed to exempt the rents associated with patent grants from business taxes, are likely also to impact the revenue collected from individual income taxes.

employees who are inventors changes appreciably in response to an allowance. We also find little evidence that the quality of new hires (“entrants”), as proxied by their earnings in the year prior to hiring (Column 3 of Table 6), rises in response to an initial patent allowance. Likewise, the earnings of those workers who choose to separate from the firm appear to be unaffected by the allowance (Column 4 of Table 6).

Examining “firm stayers” who were present in the year of application and continued to be employed by the firm provides a different window into potential changes in workforce composition. We find no appreciable effect on the application year earnings of stayers (Column 5 of Table 6), suggesting little change in the quality of retained workers. Finally, the average age of W2 employees drops by roughly a year in response to a valuable patent allowance (Column 6 of Table 6), which is in keeping with our finding that firms grow in response to valuable allowances and the fact that job mobility declines with age (Farber 1994).

Columns 7 and 8 report impacts on a pair of indices of worker “quality.” Each index gives the firm’s average in that year of the predicted log earnings of its employees. The first index forms predictions from a regression of individual log W2 earnings on a quartic in age fully interacted with gender and inventor status plus controls for tax year fixed effects (which are not used to form the prediction). The second index adds a polynomial in workers’ earnings on the previous job as a predictor along with an indicator for whether this is the worker’s first job. Impacts on both quality measures are statistically indistinguishable from zero. Taken together, these results provide no evidence of skill upgrading responses and hint that mild skill downgrading (primarily through age declines) is a more likely possibility.

6.3 Impacts on Within-Firm Inequality

Figure 7 analyzes the impact of initial allowances on various measures of within-firm inequality. The underlying estimates used to construct these figures are reported in Appendix Tables C.4 and C.5. Consistent with the literature on gender differences in rent-sharing (e.g., Black and Strahan 2001; Card, Cardoso, and Kline 2016), we find that initial allowances exacerbate the gender earnings gap. While male earnings rise by roughly \$6,000 (or roughly 10%; Column 1 of Appendix Table C.4) in response to a valuable patent allowance, female earnings appear unresponsive to initial allowances (Column 2 of Appendix Table C.4). Among firms that employ both genders, the gender earnings gap increases by roughly \$7,000 in response to a valuable initial allowance, or roughly 30% (Column 3 of Appendix Table C.4).

The earnings gap between inventors and non-inventors also widens in response to an initial allowance. Column 4 of Appendix Table C.4 shows that the earnings of inventors rise by roughly \$17,000 in response to an initial allowance. The earnings of non-inventors rise by only around \$2,000. Focusing on firms that employ both inventors and non-inventors, we find that the inventor-non inventor earnings gap increases by roughly \$15,000 in response to

a valuable initial allowance, or roughly 18% (Column 6 of Appendix Table C.4). The gender and inventor gaps are overlapping, but not identical phenomena. Figure 7 shows that the earnings of non-inventor males rise by roughly \$4,000 — less than all males, but more than all non-inventors.

Another important within-firm contrast is between firm officers and other workers. All US businesses are required to list officer pay separately from the pay of non-officers when filing taxes. Officers are employees who have the authority to delegate tasks and to hire employees for the jobs that need performing, and typically correspond to high-level management executives. We find that an initial allowance raises average officer earnings per W2 by roughly \$3,700, enough to explain nearly the entire W2 earnings response reported in Table 5. By contrast, non-officer earnings exhibit no appreciable response to initial allowances, though we cannot rule out small increases. As shown in Appendix Table C.7, the components of labor compensation other than officer earnings also fail to respond to patent allowances, suggesting that profit-sharing and employee benefit programs do not respond strongly to patents.

Finally, to provide a composite measure of within-firm earnings inequality, we break workers in each firm-year with at least four W2s into quartiles based on their annual earnings. We find no effect of an initial allowance on the average earnings of workers in the bottom three quartiles of the firm-specific earnings distribution, but the mean earnings of top-quartile workers rises by roughly \$8,000 per worker. The pay gap between top and bottom quartile workers rises by roughly the same amount (Column 10 of Appendix Table C.4).

6.4 Impacts on Earnings by Timing of Worker Entry and Exit

Our results in Section 6.2 suggested that initial allowances are not associated with major changes in workforce composition. However, an alternative way to hold constant the quality of the workforce is to study the impact of a patent allowance on the earnings of a fixed cohort of workers.

Column 1 of Table 7 documents that the average earnings of the cohort of workers present in the year of the patent application rise by roughly \$4,000 or about 7% in response to an initial allowance. These effects are concentrated in the subset of the cohort that remains with the applicant firm (“stayers”), whose earnings are estimated to rise by \$8,000 (around 11%) per year in response to an initial allowance (Column 2 of Table 7). Members of the application cohort who leave the firm, by contrast, have earnings that fall statistically insignificantly in response to an initial allowance (Column 3 of Table 7). The concentration of earnings effects on stayers casts some doubt on reputational (or “career concerns”) explanations for firm-specific wage fluctuations (Harris and Holmström 1982; Gibbons and Murphy 1992; Holmström 1999), as firm leavers appear to be unable to transport their patent-induced wage gains to new employers.

The model of Section 2 interpreted wage fluctuations as rent sharing with incumbent workers. Consistent with

that model, we find an economically small and statistically insignificant effect of initial allowances on the average earnings of entrants (Column 4 of Table 7). Given our finding in Section 6.2 that the composition of entrants does not seem to have changed in response to initial allowances, the discrepancy between our measured impacts of initial allowances on the earnings of entrants and on the earnings of firm stayers suggests that the order in which workers are hired plays an independent role in the transmission of firm shocks to wages.²³ Column 10 of Appendix Table C.3 shows that the differential earnings response of firm stayers to patent allowances is not confined to small firms.

As mentioned in Section 2, some dynamic models (e.g., Postel-Vinay and Robin 2002) can generate a drop in entry wages in response to a firm productivity increase because wage growth rates increase. Such an elevation of wage growth rates should eventually impact earnings levels. However, Column 5 of Table 7 reveals a negative (“wrong-signed”) and statistically insignificant impact of initial allowances on the earnings of workers hired within the last three years. A shift in growth rates, in conjunction with stable entry wages, should also lead to an escalating pattern of pooled wage impacts. However, we saw in Figure 4 that wage impacts are roughly stable after the initial decision. Hence, we conclude there is no evidence of a permanent impact on earnings growth rates.²⁴

Columns 6 through 8 of Table 7 adjust for possible compositional changes by subtracting from the various earnings measures an average earnings level of the same group of workers in the year of application, which adjusts for any time-invariant heterogeneity in worker quality. Column 6 (which can be compared to Column 2) shows that subtracting the average application year earnings of the firm stayers has little effect on the estimates. The estimates in Columns 7 and 8 (analogous to Columns 3 and 4) remain statistically equal to zero, suggesting that these other groups’ earnings are relatively insensitive to the patent decision.

Finally, Figure 8 reports impacts of high-value initial allowances on the average earnings of various groups of firm stayers. Initial allowances exacerbate the gender earnings gap among stayers, but the impacts on the earnings of female stayers are now estimated to be positive at around \$2,700. Appendix Table C.6 shows that the impact of an initial allowance on the earnings gap between male and female firm stayers is roughly \$9,000 (around 23%) and statistically distinguishable from zero. As one point of comparison, we find that the earnings of male firm stayers respond roughly 2.9 times as much as their female colleagues, which is slightly below the corresponding ratio of 4 found by Black and Strahan (2001) in their study of banking deregulation using firm aggregates. Likewise, earnings of inventor-stayers are estimated to increase by far more than those of non-inventors. However, the estimated impacts on non-inventor stayers are clearly distinguishable from zero and amount to a roughly 9% increase. This

²³Related work by Buhai et al. (2014) shows that worker seniority exerts an independent effect on wages even after netting out firm-wide shocks.

²⁴We have also directly computed impacts on earnings growth rates for workers hired within the last three years, but this led to highly imprecise estimates. Specifically, we estimate an impact of negative one percentage point on the three-year growth rate of the earnings of new hires, with a standard error of seven percentage points.

responsiveness of non-inventor earnings but larger response of inventor earnings echoes the findings presented in contemporaneous work by Aghion et al. (2018), which estimates that in Finnish firms inventor earnings are around twice as responsive to patent application filings as are non-inventor earnings.

The bottom of Figure 8 reports impacts on average earnings by the worker’s position in the firm’s earnings distribution *at the time of the patent application*. Large earnings gains, amounting to roughly 6-8% increases, are present for firm stayers initially in the top half of the firm-specific earnings distribution. In our estimation sample, firms with high value patents have, in an average year, roughly ten stayers in the top initial earnings quartile and nine in the third quartile. Because earnings impacts are clearly present among third quartile workers, our pooled impacts on stayer earnings are unlikely to solely represent the capture of rents by CEOs or other top executives. By contrast, the earnings response of firm-stayers initially in the bottom half of the distribution exhibit relatively muted responses, that are statistically indistinguishable from zero.

7 Pass-Through Estimates

Table 8 reports rent-sharing specifications based on equation (5) that relate earnings outcomes to surplus per worker. As discussed in Section 3.4, our preferred approach uses the sum of wages and EBITD to measure surplus. However, for comparison with past literature, we also report in Panel B specifications proxying surplus with our measure of value added.

Column 1a of Panel A shows that regressing average wage bill per worker on our preferred measure of surplus per worker, together with our standard set of (firm and art-unit by application year by calendar year) fixed effects yields an estimated pass-through coefficient $\hat{\pi}$ of 0.16. Instrumenting surplus with the interaction of a post-decision indicator and an indicator for the application being initially allowed increases the estimated coefficient to 0.29, implying that workers capture 29 cents of each additional dollar of surplus. Because our first stage F statistic is near the benchmark of 10, we also provide a weak-identification robust confidence interval, which reveals that we can reject values of π below 0.1 or above 0.57 at the 10% level.²⁵ For comparison with the prior literature, we also convert our estimates of π to elasticities using the means of surplus and wages among firms with top quintile patent applications. While OLS estimation yields a pass-through elasticity of 0.19, IV yields an estimated elasticity of 0.35. A plausible candidate explanation for the larger IV estimates is that wages respond more strongly to lower frequency fluctuations in surplus (Guiso, Pistaferri, and Schivardi 2005); however, in Appendix Table C.8 we document that using three-year averages of surplus yields very small increases in estimates of π .

²⁵These confidence intervals, which are two-way clustered on art unit and application year by decision year, employ the minimum distance variant of the Anderson-Rubin test statistic (Anderson and Rubin, 1949) described in Section 5.1 of Andrews, Stock, and Sun (2018). The endpoints of the confidence interval are defined by quadratic inequalities, which we solved analytically. We thank Liyang Sun for suggesting this approach.

Columns 1a and 1b of Panel B show the corresponding results when value added per worker is treated as the endogenous variable. This yields lower pass-through coefficients, which is in keeping with the notion that value added includes a number of extraneous cost components that cannot be captured by workers. In elasticity terms, however, using value added yields larger elasticities because value added has a greater mean than our preferred surplus measure. A useful point of comparison comes from Van Reenen (1996) who reports an elasticity of average wages with respect to quasi-rents of 0.29 (Table III, Row 2). Card et al. (2018) suggest doubling quasi-rent elasticities to make them roughly comparable to a value added elasticity. Applying this rule of thumb to Van Reenen’s study yields a value added equivalent elasticity of 0.58, which is slightly above our instrumented value added elasticity estimate of 0.47. On the other hand, our value added pass-through coefficient of 0.23 is directly comparable to the firm-level pass-through estimates of Abowd and Lemieux (1993) who report an identical pass-through coefficient of 0.23 (Table III, Col 8).

Columns 2 and 3 of Panel A change the dependent variable to be the earnings of various subgroups of workers employed by most firms.²⁶ OLS estimates indicate that the earnings of men are slightly more sensitive to surplus fluctuations than are the earnings of workers in general. However, instrumenting surplus with initial allowances dramatically raises this point estimate, indicating that men capture 53 cents of every dollar of surplus per worker, roughly 80% higher than was found for the pooled estimate. By contrast, non-inventor earnings responses are relatively muted, indicating such workers capture only 19 cents of every dollar of surplus per worker.

Column 4 of Panel A restricts attention to firm stayers, who were present in the year of application. OLS estimates indicate that stayer earnings are more sensitive to surplus fluctuations in levels (relative to the sample of all workers), but the elasticity is the same as was found for the average earnings of all workers (0.19). Instrumenting the surplus changes this conclusion dramatically: stayers are estimated to capture 61 cents of every dollar of surplus, with a corresponding elasticity of 0.56. Remarkably, the 90% confidence interval for π in this subgroup ranges from 0.21 to 1, indicating we cannot reject that firm stayers capture the entirety of their replacement costs in higher earnings. Column 5 adjusts stayer earnings for potential changes in workforce composition by subtracting off their earnings in the application year, which should difference out any selection on time invariant worker skills. As expected given our results in Section 6.2, this adjustment has minor effects on the results — lowering, for instance, the instrumented pass-through of surplus to earnings from 61 cents to 51 cents on the dollar.

Finally, Column 6 shows that our pass-through results are not driven exclusively by workers listed as inventors on patent applications: the instrumented value of π among non-inventor stayers is 0.48. Though the standard errors for Column 6b are somewhat smaller than the estimates in Column 4b, the first stage is somewhat weaker, which leads the lower limit of the 90% confidence interval for π among non-inventor stayers to be nearly identical to that of all stayers. Because most previous studies of rent-sharing do not focus on innovative firms, this estimate is

²⁶Table 8 omits estimates for subgroups (e.g. female inventors) that have sample sizes too small to produce reliable estimates.

arguably most comparable to the work reviewed in Card et al. (2018).

In sum, we find that the earnings of workers, particularly those who were present in the year of application, are quite sensitive to fluctuations in surplus. On average, a dollar increase in surplus is estimated to yield a 29 cent increase in worker earnings and a 61 cent increase in the earnings of firm stayers. Using value added instead of our preferred surplus measure yields uniformly lower pass-through estimates but tends to raise elasticities substantially. In elasticity terms, our pooled estimates are larger than the bulk of recent studies reviewed by Card et al. (2018), but align closely with the estimates of Abowd and Lemieux (1993) and Van Reenen (1996) which exploit firm aggregates.

Our finding of larger elasticities may, in part, be attributable to our use of external instruments. Abowd and Lemieux (1993), Van Reenen (1996), and Garin and Silverio (2017) all find that instrumenting value added yields large increases in rent-sharing estimates. Garin and Silverio (2017) estimate a pooled elasticity of 0.15 (Table 6, Column 4) in Portuguese data using exposure to exchange rate shocks as an instrument, and find a much larger elasticity of 0.28 (Table 9, Column 2) in industries with low separation rates. Because measuring the level of surplus is particularly difficult at small firms, we are somewhat less confident in our elasticity estimates than we are in the more theoretically motivated pass-through coefficients, which are robust to mismeasurement of the level of surplus. Nevertheless, our 90% confidence interval for π permits corresponding surplus elasticities as low as 0.19 for firm stayers.²⁷

Another plausible explanation for our finding of strong earnings sensitivity to surplus shocks is our focus on innovative firms, which are likely to rely heavily on the specific human capital of their workforce. This interpretation is consistent with the findings of Van Reenen (1996), who also studied innovative firms. Our finding of very large wage pass-through to early cohorts of workers is consistent with the notion that early employees, some of whom may be “founders,” are particularly difficult for firms to replace.

8 Retention Estimates

The wage posting model of Section 2 interpreted earnings responses to firm-specific shocks as attempts to retain incumbent workers. Figure 9 provides event study estimates of the impact of patent allowances on the logarithm of the fraction of the application cohort working at the firm, split by whether the worker was in the top or bottom half of the firm-specific earnings distribution in the year of application. Recall from Figure 8 that the earnings responses to initial allowances were concentrated in the top half of the distribution of firm stayers. Consistent with the notion that these earnings movements capture rent sharing, the retention of “above median”

²⁷This figure comes from multiplying the lower limit of the confidence interval for π , which is 0.21, by the ratio of mean stayer wages to surplus among firms with high value patent applications, which is $\frac{0.56}{0.61} \approx .92$.

firm stayers (right-hand side panel of Figure 9) responds strongly to initial allowances while the retention of “below median” stayers (left-hand side panel of Figure 9) exhibits a very weak and statistically insignificant response to allowances. Interestingly, the retention response stabilizes by three years after the initial decision date. This pattern suggests the earnings response, which from Figure 4 manifests rather quickly, serves to retain incumbent workers who would have otherwise separated over the first three years after an initial rejection.

Table 9 scales the retention responses of various groups of workers present in the application year by the impact on their log earnings to obtain IV estimates of the incumbent retention-wage elasticity. Instrumenting stayer wages with the initial allowance decision yields an estimated retention-wage elasticity of 1.2, or equivalently, a separation-wage elasticity of -1.7. This estimate is well within the range of separation elasticities reported in Manning (2011)’s review of quasi-experimental studies but is somewhat larger in magnitude than the short run elasticities reported in Dube, Giuliano, and Leonard (2018). However, Dube, Giuliano, and Leonard (2018) report nine month elasticities, while we interpret our estimates as representing three year elasticities, which we would expect to be a bit larger. Despite a first stage F statistic below the benchmark of 10, our weak-identification robust confidence interval indicates that we can reject retention elasticities below 0.44 at the 10% level.

We find little evidence of heterogeneity in the retention elasticity, though our analysis is hampered by a weak first stage for some subgroups. Among “above median” stayers, the retention elasticity rises slightly to 1.4, but we cannot reject that the elasticity is the same as in the pooled sample, which is in keeping with the notion that the pooled results are driven primarily by the above median stayers. We do not report estimates for below median stayers because the first stage is extremely weak, which leads to erratic estimates. Male retention elasticities are estimated to be somewhat below female elasticities, but the female estimates are imprecise to the point of being indistinguishable from zero. Finally, non-inventors are estimated to have a retention elasticity of 1.3, nearly identical to what we found in our pooled analysis. The finding of a stable retention elasticity across groups reinforces the evidence in Figure 9 that the groups experiencing the largest earnings responses also exhibit the largest retention responses. This corroborates our model-based interpretation of the earnings impacts we measure as reflecting economic rents, a view we consider in more quantitative detail in Section 9.

9 Model-based interpretation

In the model of Section 2, the retention wage elasticity can be written:

$$\frac{d \ln G(w_j^I)}{d \ln w_j^I} = \eta \frac{w_j^I}{w_j^I - w_j^m} = \eta \frac{w_j^I / w_j^m}{w_j^I / w_j^m - 1}.$$

Hence, we require a calibration of w_j^l/w_j^m to recover η from the estimates in Table 9. From Table 2, workers hired within the three years prior to the year of application earn roughly \$43,500, which we take as a measure of the entry wage w_j^m . By contrast, workers who have been at the firm for four or more years earn roughly \$79,000, which we take as a measure of w_j^l . Hence, we calibrate $w_j^l/w_j^m = 79/43.5 \approx 1.8$.

In Table 9, we found a pooled retention elasticity of approximately 1.2. Hence, our estimate of η is $1.2 \times \frac{1.8}{0.8} \approx 2.7$. Recall from equation (3) that in the model of section 2 workers are offered a fraction $\theta = \frac{\eta}{1+\eta}$ of their marginal replacement costs as a wage premium. Our retention elasticity estimate therefore implies that incumbent workers capture roughly 73% of their replacement costs in wage premia.

We can also use our estimates to quantify these marginal replacement costs. Rearranging equation 3, we have $\frac{c'(N_j/I_j)}{w_j^m} = [w_j^l/w_j^m - 1] / \theta = .8/.73 \approx 1.1$. Hence, our calibration suggests that the marginal replacement cost of an incumbent worker is roughly equal to the annual earnings of a new hire. This replacement cost estimate is higher than is usually found in simple linear-quadratic models of employment adjustment (Hamermesh and Pfann 1996; Bloom 2009; Cooper and Willis 2009). However, we study fairly large shocks to small firms which, with convexity in hiring / training costs, should lead to correspondingly large replacement costs on the margin.

We can also use our estimates to compute an implied elasticity of product demand ε . In Table 8 we found that incumbent workers captured 61 cents of every dollar of patent induced surplus. Taking $\pi = .61 = (\frac{\varepsilon-1}{\varepsilon}) \theta$ and using our estimate of $\theta = .73$ implies that $\varepsilon \approx 6.0$, which corresponds to a 20% markup of product price over marginal cost. This finding is in line with recent work that has used values of ε ranging from 4.5 (Suárez Serrato and Zidar 2016) to 7 (Coibion, Gorodnichenko, and Wieland 2012).

Appendix Table C.9 reports some alternative calibrations of model inputs that set the pass-through and retention elasticities to different values along with the incumbent wage premium. Interestingly, some calibrations yield invalid values of the structural parameters, suggesting our model can be used to rule out some configurations of parameters falling within our confidence intervals. The general theme of this sensitivity exercise is that, across a wide range of potential rationalizations of the data, workers capture a large fraction of their marginal replacement costs in wage premia and that those costs are substantial.

It is worth remarking briefly on how our model rationalizes the gender differences in earnings pass-through reported in Figure 8. The model suggests two possible explanations for these differences. A first potential explanation is that men and women might face different distributions of outside offers, which would manifest in different retention elasticities and consequently different pass-through coefficients. However, the results of Table 9 provide little support for this conjecture. If anything, women exhibit slightly higher retention elasticities than men, which should yield greater earnings pass-through for them.

A second potential explanation for gender differences in earnings pass-through is that the marginal replacement costs of men could — on average — exceed those of women. Concentration of women in occupations involving

smaller training and recruiting costs, for example, could plausibly generate such differences. Recall that earnings impacts are concentrated among firm “officers” who are likely difficult to replace because of the specific capital embedded in their relationships with subordinate workers. Unfortunately, because of the way in which officer earnings are reported (as aggregates in the firm-level data, rather than as a variable in our worker micro-data) we do not know what fraction of officers are women. However, for the average firm in our sample, the fraction of women in the top quartile of earnings distribution in the year of initial patent application is only 11.5%, a fact that is consistent with a broad range of evidence suggesting that US women tend to be employed in lower paying occupations than men (Goldin 2014).

10 Conclusion

This paper analyzes how patent-induced shocks to labor productivity propagate into worker earnings using a new linkage of US patent applications to US business and worker tax records. Our baseline estimates suggest that on average every patent-induced dollar of surplus yields roughly 30 cents of additional earnings; this share is roughly twice as high for incumbent workers present since the year of application. Among non-inventors present since the year of application, who are arguably the group most comparable to the recent studies reviewed by Card et al. (2018), we find a both a pass-through rate and elasticity of roughly a half. These estimates provide some of the first evidence, along with Jäger (2016), that truly idiosyncratic variability in firm performance is an important causal determinant of worker pay. Given that firm productivity is highly variable and persistent (Luttmer 2007; Foster, Haltiwanger, and Syverson 2008), it is plausible that firm-specific shocks contribute substantially to permanent earnings inequality among identically skilled workers.

We document several sources of heterogeneity in the pass-through of patent-induced shocks to workers. First, patent allowances have no effect on the earnings of new hires. This finding may be specific to the small firms we study, which are unlikely to exhibit market power over new hires. Nevertheless, this finding implies that patent shocks “stretch” the firm’s pay scale by increasing inequality between new hires and incumbent workers. Second, among incumbent workers, patent allowances exacerbate the within-firm gender earnings gap. The gender differences in earnings pass-through found here are larger than those estimated by Card, Cardoso, and Kline (2016) and Garin and Silverio (2017) in Portuguese data, but smaller than those reported by Black and Strahan (2001) in US data. Third, while the earnings of both inventors and non-inventors respond to patent decisions, the earnings of inventors are substantially more responsive, which is notable because previous studies of pass-through to inventors have studied settings where inventor compensation is mandated by law.²⁸ Finally, earnings impacts are strongly

²⁸For example, Aghion et al. (2018) analyze how inventor and non-inventor earnings change before and after patent applications among Finnish firms, but Finland — like many other European countries — has a law that requires firms to pay inventors for inventions produced while they are employed. See the discussion in Toivanen and Väänänen (2012).

concentrated among employees in the top quartile of the within-firm earnings distribution, among firm “officers,” and among firm stayers initially in the top half of the earnings distribution.

Two aspects of these heterogeneous earnings estimates are worth emphasizing. First, these impacts appear to mirror heterogeneity in the costs of replacing different types of workers. Substituting new hires for high-skilled incumbents is particularly difficult. Our retention results corroborate this view: worker retention rises most strongly among groups of workers with the largest earnings increases. This pattern suggests, via revealed preference, that these earnings fluctuations constitute economic rents. A quantification of our model finds that incumbent workers capture the majority of their replacement costs in wage premia. The pairing of incumbent rents of this magnitude with stable new hire earnings highlights the importance of seniority and specific investments in wage determination — themes emphasized by, among others, Becker (1964), Stevens (1994), and Manning (2006). Second, our findings strongly suggest that firm shocks play an important role in generating earnings inequality not only *across* but also *within* workplaces. Understanding the extent to which heterogeneity in pass-through across workers contributes to overall earnings inequality is an important topic for future research.

References

- Abowd, John and Thomas Lemieux**, “The Effects of Product Market Competition on Collective Bargaining Agreements: The Case of Foreign Competition in Canada,” *Quarterly Journal of Economics*, 1993, 108 (4), 983–1014.
- , **Francis Kramarz, and David Margolis**, “High Wage Workers and High Wage Firms,” *Econometrica*, 1999, 67 (2), 251–333.
- , **Kevin McKinney, and Nellie Zhao**, “Earnings Inequality and Mobility Trends in the United States: Nationally Representative Estimates from Longitudinally Linked Employer-Employee Data,” *Journal of Labor Economics*, 2018, 26 (S1), S183–S300.
- Acemoglu, Daron and William Hawkins**, “Search with Multi-Worker Firms,” *Theoretical Economics*, 2014, 9 (3), 583–628.
- Aghion, Philippe, Ufuk Akcigit, Ari Hyytinen, and Otto Toivanen**, “On the returns to invention within firms: Evidence from Finland,” *AEA Papers and Proceedings*, 2018, 108, 208–12.
- Anderson, Theodore W and Herman Rubin**, “Estimation of the parameters of a single equation in a complete system of stochastic equations,” *The Annals of Mathematical Statistics*, 1949, 20 (1), 46–63.
- Andrews, Isaiah, James Stock, and Liyang Sun**, “Weak Instruments in IV Regression: Theory and Practice,” Technical Report 2018. MIT Working Paper.
- Axtell, Robert L**, “Zipf distribution of US firm sizes,” *Science*, 2001, 293 (5536), 1818–1820.
- Balasubramanian, Natarajan and Jagadeesh Sivadasan**, “What Happens When Firms Patent? New Evidence from U.S. Economic Census Data,” *Review of Economics and Statistics*, 2011, 93 (1), 126–146.
- Barth, Erling, Alex Bryson, James Davis, and Richard Freeman**, “It’s Where You Work: Increases in Earnings Dispersion across Establishments and Individuals in the U.S.,” *Journal of Labor Economics*, 2016, 34 (S2), S67–S97.
- Becker, Gary S**, *Human Capital Theory*, New York, NY: Columbia University Press, 1964.
- Bell, Alex, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen**, “The Lifecycle of Inventors,” 2016. Working paper.
- Bender, Stefan, Nicholas Bloom, David Card, John Van Reenen, and Stefanie Wolter**, “Management Practices, Workforce Selection and Productivity,” *Journal of Labor Economics*, 2018, 36 (S1), S371–S409.
- Bessen, James**, “The Value of U.S. Patents by Owner and Patent Characteristics,” *Research Policy*, 2008, 37 (5), 932–945.
- Black, Sandra and Philip Strahan**, “The Division of the Spoils: Rent-Sharing and Discrimination in a Regulated Industry,” *American Economic Review*, 2001, 91 (4), 814–831.
- Blanchflower, David, Andrew Oswald, and Peter Sanfey**, “Wages, Profits, and Rent-Sharing,” *Quarterly Journal of Economics*, 1996, 111 (1), 227–251.
- Bloom, Nicholas**, “The Impact of Uncertainty Shocks,” *Econometrica*, 2009, 77 (3), 623–685.
- **and John Van Reenen**, “Measuring and Explaining Management Practices Across Firms and Countries,” *Quarterly Journal of Economics*, 2007, 122 (4), 1351–1408.


- , **Mark Schankerman**, and **John Van Reenen**, “Identifying Technology Spillovers and Product Market Rivalry,” *Econometrica*, 2013, 81 (4), 1347–1393.
- Brown, James and Orley Ashenfelter**, “Testing the Efficiency of Employment Contracts,” *Journal of Political Economy*, 1986, 94 (3), S40–S87.
- Buhai, I. Sebastian, Miguel A. Portela, Coen N. Teulings, and Aico van Vuuren**, “Returns to Tenure or Seniority,” *Econometrica*, 2014, 82 (2), 705–30.
- Cahuc, Pierre and Etienne Wasmer**, “Does Intrafirm Bargaining Matter in the Large Firm’s Matching Model?,” *Macroeconomic Dynamics*, 2001, 5 (5), 742–747.
- Card, David, Ana Rute Cardoso, and Patrick Kline**, “Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women,” *Quarterly Journal of Economics*, 2016, 131 (2), 633–686.
- , —, **Jörg Heining, and Patrick Kline**, “Firms and Labor Market Inequality: Evidence and Some Theory,” *Journal of Labor Economics*, 2018, 36 (S1), S13–S70.
- , **Francesco Devicienti, and Agata Maida**, “Rent-Sharing, Holdup, and Wages: Evidence from Matched Panel Data,” *Review of Economic Studies*, 2014, 81 (1), 84–111.
- , **Jörg Heining, and Patrick Kline**, “Workplace Heterogeneity and the Rise of West German Wage Inequality,” *Quarterly Journal of Economics*, 2013, 128 (3), 967–1015.
- Carlsson, Mikael, Julián Messina, and Oskar Nordstrom Skans**, “Wage Adjustment and Productivity Shocks,” *Economic Journal*, 2016, 126 (595), 1739–1773.
- Christofides, Louis and Andrew Oswald**, “Real Wage Determination and Rent-Sharing in Collective Bargaining Agreements,” *Quarterly Journal of Economics*, 1992, 107 (3), 985–1002.
- Cohen, William, Pradeep Ravikumar, and Stephen Fienburg**, “A Comparison of String Distance Metrics for Name-Matching Tasks,” in “IIWeb” 2003, pp. 73–78.
- Coibion, Olivier, Yuriy Gorodnichenko, and Johannes Wieland**, “The Optimal Inflation Rate in New Keynesian Models: Should Central Banks Raise Their Inflation Targets in Light of the Zero Lower Bound?,” *Review of Economic Studies*, 2012, 79 (4), 1371–1406.
- Cooper, Russell and Jonathan L Willis**, “The Cost of Labor Adjustment: Inferences from the Gap,” *Review of Economic Dynamics*, 2009, 12 (4), 632–647.
- DeBacker, Jason Matthew, Bradley Heim, Shanthi Ramnath, and Justin Ross**, “The Impact of State Taxes on Pass-Through Businesses: Evidence from the 2012 Kansas Income Tax Reform,” 2016. Working paper.
- Depalo, Domenico and Sabrina Di Addario**, “Shedding light on inventors’ returns to patents,” Technical Report 2014. IRL Working Paper #115-14.
- Dube, Arindrajit, Laura Giuliano, and Jonathan Leonard**, “Fairness and frictions: The impact of unequal raises on quit behavior,” Technical Report, National Bureau of Economic Research 2018.
- Farber, Henry**, “The Analysis of Interfirm Worker Mobility,” *Journal of Labor Economics*, 1994, 12 (4), 554–593.
- Farre-Mensa, Joan, Deepak Hegde, and Alexander Ljungqvist**, “What Is a Patent Worth? Evidence from the U.S. Patent ‘Lottery,’” 2017. NBER working paper no. 23268.

- Feenberg, Daniel and Elisabeth Coutts**, “An Introduction to the TAXSIM Model,” *Journal of Policy Analysis and Management*, 1993, 12 (1), 189–194.
- Feigenbaum, James**, “Automated Census Record Linking: A Machine Learning Approach,” 2016. Working paper.
- Foster, Lucia, John Haltiwanger, and Chad Syverson**, “Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?,” *American Economic Review*, 2008, 98 (1), 394–425.
- Frakes, Michael and Melissa Wasserman**, “Is the Time Allocated to Review Patent Applications Inducing Examiners to Grant Invalid Patents? Evidence from Micro-Level Application Data,” *Review of Economics and Statistics*, 2017, 99 (3), 550–563.
- Friedrich, Benjamin, Lisa Laun, Costas Meghir, and Luigi Pistaferri**, “Earnings dynamics and firm-level shocks,” 2014.
- Garin, Andrew and Filipe Silverio**, “How Does Firm Performance Affect Wages? Evidence from Idiosyncratic Export Shocks,” Technical Report, Working Paper 2017.
- Gibbons, Robert and Kevin J. Murphy**, “Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence,” *Journal of Political Economy*, 1992, 100 (3), 468–505.
- Goldberg, Pinelopi Koujianou and Rebecca Hellerstein**, “A Structural Approach to Identifying the Sources of Local Currency Price Stability,” *Review of Economic Studies*, 2013, 80 (1), 175–210.
- Goldin, Claudia**, “A Grand Gender Convergence: Its Last Chapter,” *American Economic Review*, 2014, 104 (4), 1091–1119.
- Goldschmidt, Deborah and Johannes Schmieder**, “The Rise of Domestic Outsourcing and the Evolution of the German Wage Structure,” *Quarterly Journal of Economics*, 2017, 132 (3), 1165–1217.
- Gorodnichenko, Yuriy and Oleksandr Talavera**, “Price Setting in Online Markets: Basic Facts, International Comparisons, and Cross-Border Integration,” *American Economic Review*, 2017, 107 (1), 249–282.
- Graham, Stuart and Deepak Hegde**, “Do Inventors Value Secrecy in Patenting? Evidence from the American Inventor’s Protection Act of 1999,” 2014. Unpublished USPTO mimeo.
- Guiso, Luigi, Luigi Pistaferri, and Fabiano Schivardi**, “Insurance Within the Firm,” *Journal of Political Economy*, 2005, 113 (5), 1054–1087.
- Hall, Bronwyn, Adam Jaffe, and Manuel Trajtenberg**, “The NBER U.S. Patent Citations Data File: Lessons, Insights, and Methodological Tools,” 2001. NBER working paper no. 8498.
- Hall, Robert and Paul Milgrom**, “The Limited Influence of Unemployment on the Wage Bargain,” *American Economic Review*, 2008, 98 (4), 1653–1674.
- Hamermesh, Daniel S and Gerard A Pfann**, “Adjustment costs in factor demand,” *Journal of Economic Literature*, 1996, 34 (3), 1264–1292.
- Harhoff, Dietmar, Frederic Scherer, and Katrin Vopel**, “Citations, Family Size, Opposition, and the Value of Patent Rights,” *Research Policy*, 2003, 32 (8), 1343–1363.
- Harrell, Frank**, *Regression Modeling Strategies: With Applications to Linear Models, Logistic Regression, and Survival Analysis*, corrected ed., New York, NY: Springer, 2001.
- Harris, Milton and Bengt Holmström**, “A Theory of Wage Dynamics,” *The Review of Economic Studies*, 1982, 49 (3), 315–333.

- Hastie, Trevor, Robert Tibshirani, and Jerome Friedman**, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed., New York, NY: Springer, 2016.
- Helpman, Elhanan, Oleg Itskhoki, Marc-Andreas Muendler, and Stephen Redding**, “Trade and Inequality: From Theory to Estimation,” *Review of Economic Studies*, 2017, 84 (1), 357–405.
- Hildreth, Andrew**, “Rent-Sharing and Wages: Product Demand or Technology Driven Premia,” *Economics of Innovation and New Technology*, 1998, 5 (2-4), 199–226.
- Holmström, Bengt**, “Moral Hazard and Observability,” *The Bell Journal of Economics*, 1979, 10 (1), 74–91.
- , “Agency Costs and Innovation,” *Journal of Economic Behavior & Organization*, 1989, 12 (3), 305–327.
- , “Managerial Incentive Problems: A Dynamic Perspective,” *The Review of Economic Studies*, 1999, 66 (1), 169–182.
- Jaffe, Adam, Manuel Trajtenberg, and Rebecca Henderson**, “Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations,” *Quarterly Journal of Economics*, 1993, 108 (3), 577–598.
- Jäger, Simon**, “How Substitutable Are Workers? Evidence from Worker Deaths,” Technical Report 2016.
- Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman**, “Technological Innovation, Resource Allocation, and Growth,” *Quarterly Journal of Economics*, 2017, 132 (2), 665–712.
- Lamadon, Thibaut**, “Productivity Shocks, Long-Term Contracts and Earnings Dynamics,” 2016. Working paper.
- Lanjouw, Jean and Mark Schankerman**, “Characteristics of Patent Litigation: A Window on Competition,” *RAND Journal of Economics*, 2001, 32 (1), 129–151.
- Lemley, Mark and Bhaven Sampat**, “Is the Patent Office a Rubber Stamp?,” *Emory Law Journal*, 2008, 58, 181–203.
- and —, “Examining Patent Examination,” *Stanford Technology Law Review*, 2010, (4), 1–11.
- Lentz, Rasmus and Dale Mortensen**, “Labor Market Models of Worker and Firm Heterogeneity,” *Annual Review of Economics*, 2010, 2, 577–602.
- Luttmer, Erzo**, “Selection, Growth, and the Size Distribution of Firms,” *Quarterly Journal of Economics*, 2007, 122 (3), 1103–1144.
- Manning, Alan**, “A Generalised Model of Monopsony,” *Economic Journal*, 2006, 116 (508), 84–100.
- , “Imperfect Competition in the Labor Market,” in Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics*, Vol. 4B, North Holland, 2011, pp. 973–1041.
- Mogstad, Magne, Bradley Setzler, Thibaut Lamadon et al.**, “Earnings Dynamics, Mobility Costs, and Transmission of Market-Level Shocks,” 2017.
- Pakes, Ariel**, “Patents as Options: Some Estimates of the Value of Holding European Patent Stocks,” *Econometrica*, 1986, 54 (4), 755–784.
- Pissarides, Christopher**, *Equilibrium Unemployment Theory*, 2nd ed., Cambridge, Mass: The MIT Press, 2000.
- , “The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?,” *Econometrica*, 2009, 77 (5), 1339–1369.


- Postel-Vinay, Fabien and Jean-Marc Robin**, “Equilibrium Wage Dispersion with Worker and Employer Heterogeneity,” *Econometrica*, 2002, 70 (6), 2295–2350.
- Putnam, Jonathan**, “The Value of International Patent Rights,” 1996. Yale PhD dissertation.
- Rosen, Sherwin**, “The Theory of Equalizing Differences,” in Orley Ashenfelter and Richard Layard, eds., *Handbook of Labor Economics*, Vol. 1, North Holland, 1986, pp. 641–692.
- Sampat, Bhaven and Heidi Williams**, “How Do Patents Affect Follow-on Innovation? Evidence from the Human Genome,” *American Economic Review*, forthcoming.
- Schankerman, Mark and Ariel Pakes**, “Estimates of the Value of Patent Rights in European Countries During the Post-1950 Period,” *Economic Journal*, 1986, 96 (384), 1052–1076.
- Serrano, Carlos**, “The Dynamics of the Transfer and Renewal of Patents,” *RAND Journal of Economics*, 2010, 41 (4), 686–708.
- Smith, Matthew, Danny Yagan, Owen Zidar, and Eric Zwick**, “Capitalists in the Twenty-First Century,” 2017. Working paper.
- Song, Jae, David Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter**, “Firming Up Inequality,” *Quarterly Journal of Economics*, forthcoming.
- Sorkin, Isaac**, “Ranking firms using revealed preference,” *Quarterly Journal of Economics*, 2018, 133 (3), 1331–1393.
- Stevens, Margaret**, “A theoretical model of on-the-job training with imperfect competition,” *Oxford Economic Papers*, 1994, pp. 537–562.
- Stole, Lars and Jeffrey Zwiebel**, “Intra-Firm Bargaining under Non-Binding Contracts,” *The Review of Economic Studies*, 1996, 63 (3), 375–410.
- Suárez Serrato, Juan Carlos and Owen Zidar**, “Who Benefits from State Corporate Tax Cuts? A Local Labor Markets Approach with Heterogeneous Firms,” *American Economic Review*, 2016, 106 (9), 2582–2624.
- Toivanen, Otto and Lotta Väänänen**, “Returns to Inventors,” *Review of Economics and Statistics*, 2012, 94 (4), 1173–1190.
- Trajtenberg, Manuel**, *Economic Analysis of Product Innovation: The Case of CT Scanners*, Cambridge, MA: Harvard University Press, 1990.
- US Bureau of Economic Analysis**, “National Income and Product Accounts Table 1.1.4: Price Indexes for Gross Domestic Product,” 2014.
- Van Reenen, John**, “The Creation and Capture of Rents: Wages and Innovation in a Panel of U.K. Companies,” *Quarterly Journal of Economics*, 1996, 111 (1), 195–226.
- Weyl, Glen and Michal Fabinger**, “Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition,” *Journal of Political Economy*, 2013, 121 (3), 528–583.
- Williams, Heidi**, “How Do Patents Affect Research Investments?,” *Annual Review of Economics*, 2017, 9, 441–469.
- Wooldridge, Jeffrey**, *Econometric Analysis of Cross Section and Panel Data*, 2nd ed., Cambridge, MA: MIT Press, 2010.
- Yagan, Danny**, “Capital Tax Reform and the Real Economy: The Effects of the 2003 Dividend Tax Cut,” *American Economic Review*, 2015, 105 (12), 3531–3563.

Figure 1: KPSS Value (ξ): Predicted Versus Actual




Notes: This figure is a binned scatterplot of actual versus predicted values of the KPSS measure of patent value ξ in millions of 1982 dollars. The sample is the subset of patent applications with non-missing values for the KPSS measure of patent value ξ . Predictions are formed based on estimates from the random effects Poisson model described in Section 5. The data in this figure have been grouped into twenty equal-sized bins. In the microdata, the slope is 1.12, as reported in the text. Here, the coefficient β instead reports the two-stage least squares slope using twenty bin dummies as instruments for predicted values and “se” reports the associated standard error.

Figure 2: Impacts by Predicted Patent Value: Surplus and Wage Bill




Notes: This figure shows the impact of an initial patent allowance on surplus per worker and wage bill per worker as a function of predicted patent value in our analysis sample. The vertical, red line is the cut-off value for the top-quintile predicted patent value subsample, and is equal to 5.3 million in 1982 USD. Values along the x-axis for the surplus series are offset from their integer value to improve readability. Surplus is EBITD (earnings before interest, tax, and depreciation) + W2 wage bill. 95% confidence intervals shown based upon standard errors two-way clustered by (1) art unit, and (2) application year by decision year.

Figure 3: Event Study Estimates: Surplus (EBITD+Wage Bill) per Worker




Notes: This figure plots the response of surplus per worker following an initial patent allowance, separately for high and low ex-ante valuable patent applications, in our analysis sample. Regressions include art unit by application year by calendar year fixed effects and firm fixed effects, as in equation (7). Standard errors are two-way clustered by (1) art unit, and (2) application year by decision year. The horizontal short-dashed, red line is the pooled differences-in-differences estimate of the impact of winning a valuable patent on surplus per worker from Table 5. Surplus is EBITD (earnings before interest, tax, and depreciation) + wage bill. $Q5$ is quintile five of predicted patent value. $<Q5$ are the remaining four quintiles. 95% confidence intervals shown. $Q5$ coefficients are offset from their integer x-axis value to improve readability.

Figure 4: Event Study Estimates: Wage Bill per Worker




Notes: This figure plots the response of wage bill per worker following an initial patent allowance, separately for high and low ex-ante valuable patent applications, in our analysis sample. Regressions include art unit by application year by calendar year fixed effects and firm fixed effects, as in equation (7). Standard errors are two-way clustered by (1) art unit, and (2) application year by decision year. The horizontal short-dashed, red line is the pooled differences-in-differences estimate of the impact of winning a valuable patent on wage bill per worker from Table 5. $Q5$ is quintile five of predicted patent value. $< Q5$ are the remaining four quintiles. $Q5$ coefficients are offset from their integer x-axis value to improve readability. 95% confidence intervals shown.

Figure 5: Event Study Estimates: $\log(\text{employees})$




Notes: This figure plots the response of the logarithm of employees per worker following an initial patent allowance, separately for high and low ex-ante valuable patent applications, in our analysis sample. Regressions include art unit by application year by calendar year fixed effects and firm fixed effects, as in equation (7). Standard errors are two-way clustered by (1) art unit, and (2) application year by decision year. The horizontal short-dashed, red line is the pooled differences-in-differences estimate of the impact of winning a valuable patent on the logarithm of the number of employees at the firm in thousands of people. $Q5$ is quintile five of predicted patent value. $< Q5$ are the remaining four quintiles. 95% confidence intervals shown. Values along the x-axis for the difference in $Q5$ are offset from their integer value to improve readability.

Figure 6: Event Study Estimates: Probability of Patent Grant




Notes: This figure plots the response of the probability of patent grant following an initial patent allowance, separately for high and low ex-ante valuable patent applications, in our analysis sample. Regressions include art unit by application year by calendar year fixed effects and firm fixed effects, as in equation (7). Standard errors are two-way clustered by (1) art unit, and (2) application year by decision year. The horizontal short-dashed, red line is the pooled difference-in-differences estimate of the impact of winning a valuable patent on the probability of the patent having been granted. Values along the x-axis for the difference in $Q5$ are offset from their integer value to improve readability. $Q5$ is quintile five of predicted patent value. $<Q5$ are the remaining four quintiles. 95% confidence intervals shown.

Figure 7: Within-Firm Inequality




Notes: This figure reports difference-in-differences coefficient and percent impact estimates of the effect of initial patent allowances on within-firm inequality measures, for high ex-ante valuable patent applications, in our analysis sample. Point estimates in the "Coefficients" panel correspond to coefficients on interactions of the designated value category with a post-decision indicator and an indicator for the application being initially allowed. Controls include main effect of value category interacted with a post-decision indicator, firm fixed effects, and art unit by application year by calendar year fixed effects, as in equation (8). 95% confidence intervals were constructed from standard errors two-way clustered by (1) art unit, and (2) application year by decision year. "Percent Impacts" point estimates correspond to the percent change in the outcome variable at the outcome variable's mean for winning a patent allowance for a high ex-ante valuable patent application. Some confidence intervals were truncated to ease visualization. Officers earnings are derived from each firm's tax filings, where firms are required to list officer and non-officer pay. Quartiles refer to within-firm wage quartiles (e.g., "Q1 earnings" measures the average wage bill in within-firm wage quartile one). Earnings are measured in thousands of 2014 USD.

Figure 8: Within-Firm Inequality Among Stayers



Notes: This figure reports difference-in-differences coefficient and percent impact estimates of the effect of initial patent allowances on within-firm stayer inequality measures, for high ex-ante valuable patent applications, in our analysis sample. Stayers are defined as those who were employed by the same firm in the year of application. Point estimates in the "Coefficients" panel correspond to coefficients on interactions of the designated value category with a post-decision indicator and an indicator for the application being initially allowed. Controls include main effect of value category interacted with a post-decision indicator, firm fixed effects, and art unit by application year by calendar year fixed effects, as in equation (8). 95% confidence intervals were constructed from standard errors two-way clustered by (1) art unit, and (2) application year by decision year. Some confidence intervals were truncated to ease visualization. Officers earnings are derived from each firm's tax filings, where firms are required to list officer and non-officer pay. Quartiles refer to within-firm wage quartiles (e.g., "Q1 earnings" measures the average wage bill in within-firm wage quartile one). Earnings are measured in thousands of 2014 USD.

Figure 9: Event Study Estimates: Employee Retention Rate by Application Year Earnings



Notes: This figure plots the response of log retention rate following an initial allowance, separately for high and low ex-ante valuable patent applications, in our analysis sample. Regressions include art unit by application year fixed effects and firm fixed effects, as in equation (7). Standard errors are two-way clustered by (1) art unit, and (2) application year by decision year. “Below median wage workers” and “Above median wage workers” respectively refer to members of the application cohort who earned below and above that firm’s median in the application year. Values along the x-axis for the difference in $Q5$ are offset from their integer value to improve readability. $Q5$ is quintile five of predicted patent value. $< Q5$ are the remaining four quintiles. 95% confidence intervals shown.

Table 1: Sample Construction

	Application- assignee pairs	Applications	Assignees	EINs
Panel A: USPTO sample				
Full sample	3,737,351	3,601,913	317,370	—
Filed between 2000 and 2010	3,063,980	2,954,507	279,936	—
Non-missing assignees	2,708,829	2,599,373	279,935	—
Non-child applications	1,341,843	1,295,649	130,619	—
Utility applications	1,339,146	1,293,054	130,113	—
First application by assignee	130,113	125,018	130,113	—
No prior grant to assignee	99,871	95,767	99,871	—
Panel B: USPTO-tax merge				
	—	39,452	39,814	81,934
First application by EIN	—	37,714	—	81,877
No prior grant to EIN	—	35,643	—	78,291
EIN with largest revenue	—	35,643	—	35,643
Active firms	—	9,732	—	9,732

Notes: This table describes the construction of our analysis sample. When selecting the first application by each assignee by date of filing (“First application by assignee”), ties are broken by taking the smallest application number. When selecting the first application for each EIN (“First application by EIN”), we drop EINs with more than one first application. When removing assignees (“No prior grant to assignee”) and EINs (“No prior grant by EIN”) with prior grants, we do so by checking against the assignees and EINs for the census of patents granted since 1976 and filed before 29 November 2000. When selecting the EIN with the largest revenue (“EIN with largest revenue”), we compare based on the revenue in the year of the application. Active firms are defined as EINs with non-zero/non-missing total income or total deductions in the application year and in the three previous years, a positive number of employees in the application year, and revenue less than 100 million in 2014 USD.

Table 2: Summary Statistics

	1. Analysis sample					2. Top quintile dosage sample				
	Mean	p10	p50	p90		Mean	p10	p50	p90	
Panel A: firm outcomes										
Revenue	9,841	226.35	3,232	29,082		12,945	353.01	4,891	39,614	
Value added	3,952	117.63	1,300	10,456		5,289	173.36	1,955	14,430	
EBITD	104.05	-916.47	87.27	1,979		12.88	-1,837	110.46	2,741	
Employment	45.65	2	17.07	117.17		61.49	2.9	25.26	152.37	
Value added per worker	116.39	15.66	84.72	287.36		120.16	14.46	86.87	315.3	
EBITD per worker	0.82	-70.88	6.85	69.53		0.72	-75.86	6.93	77.64	
Predicted patent value	4,921	219	1,749	10,942		18,104	6,089	10,502	36,651	
% Patents initially allowed	8.4	.	.	.		7.7	.	.	.	
Panel B: worker outcomes										
Labor compensation	2,132	67.88	703.44	5,410		3,088	106.7	1,111	8,310	
Wage bill	2,412	71.4	840.86	6,487		3,375	116.08	1,361	9,425	
Labor compensation per worker	57.22	10.6	43.12	136.34		60.74	12.02	47.46	148.32	
Wage bill per worker	54.98	17.94	47.94	109.33		58.42	22.35	51.55	115.04	
Average W2 earnings (<4 yrs at firm)	43.55	12.39	35.11	94.44		46.46	14.41	37.48	101.09	
Average W2 earnings (4+ yrs at firm)	78.98	29.66	63.81	147.87		82.16	35.14	66.22	150.94	
% Female employment	30.2	0	25.6	66.1		29.8	0	26.3	62.6	
% Contractors	18.4	0	10.5	53.0		17.8	0	11.0	50.3	
% Entrants	28.8	0	24.6	62.9		29.9	0	25.3	62.9	
% Inventors	11.4	0	2.5	36.4		10.5	0	3.3	32.3	
firm observations			9,732					1,946		

Notes: This table tabulates summary statistics on firm (Panel A) and worker (Panel B) outcomes for each of two samples: our analysis sample of matched patent application-firm pairs, and our sub-sample of patent application-firm pairs in the analysis sample for which the patent applications are in the top quintile of predicted value (“dosage”). All summary statistics are measured in the year in which the patent application was filed. For each variable in each panel for each sample, we tabulate the mean, 10th percentile (median), and 90th percentile; to protect taxpayer anonymity, p10, p50, and p90 refer to centile means. Entrants are defined as those employees who were not employed at the firm in the previous year. EBITD is earnings before interest, taxes, and deductions. Revenue, value added, EBITD, labor compensation, wage bill, entrant wages and incumbent wages are reported in thousands of 2014 USD. Predicted patent value is reported in thousands of 1982 USD. All percent or “per worker” outcomes in Panel B are per W2 employee except for “% Contractors” which is per W2 employee + 1099 independent contractor. For all other variable definitions, please see the text of the paper and Appendix A.

Table 3: Balance of Assignee Characteristics Across Initially Allowed and Initially Rejected Patent Applications

	Initially allowed			
	Analysis sample		Top quintile	
	(1a)	(1b)	(2a)	(2b)
log(employees)	-3.71 (1.85)	-2.06 (2.18)	-0.16 (4.70)	1.97 (4.76)
Revenue per worker	0.03 (0.01)	0.02 (0.02)	0.01 (0.02)	0.01 (0.03)
Value added per worker	-0.14 (0.05)	-0.07 (0.06)	-0.07 (0.08)	0.00 (0.12)
Wage bill per worker	0.14 (0.10)	0.14 (0.13)	0.08 (0.21)	-0.11 (0.24)
EBITD per worker	0.11 (0.05)	0.06 (0.07)	0.17 (0.10)	0.12 (0.14)
Observations	9,732	8,647	1,946	1,666
AU-AY FEs		✓		✓
p-value	0.005	0.494	0.518	0.830

Notes: This table reports covariate balance tests for initial patent allowances. Specifically, the coefficients report linear probability model estimates of the marginal effect of the included covariate on the probability that a patent application receives an initial allowance; all coefficients have been multiplied by 1,000 for ease of interpretation. AU-AY FEs denotes the inclusion of art unit (AU) by application year (AY) fixed effects. Covariates are measured as of the year of application. Columns (1a) and (1b) report the results for observations in the analysis sample. Columns (2a) and (2b) report the results for observations in the top-quintile predicted patent value. Singleton observations are dropped in the fixed effects specifications, which accounts for the smaller number of observations in Column (1b) relative to Column (1a) and in Column (2b) relative to Column (2a). Standard errors (reported in parentheses) are two-way clustered by art unit and application year by decision year except in Column (2b) which clusters by art unit (because the estimated two-way variance covariance matrix was singular). The p-value reports the probability that the covariates measured in the year of application do not influence the probability of an initial allowance. EBITD is earnings before interest, taxes, and deductions. Revenue, value added, wage bill, and EBITD are measured in thousands of 2014 USD.

Table 4: Prediction of KPSS Patent Value Based on Patent Application and Assignee Characteristics

	KPSS value (ξ)	
$\mathbf{1}(\text{patent family size} = 1)$	0.28	(0.06)
$\log(\text{patent family size})$	0.23	(0.04)
$\mathbf{1}(\text{number of claims} = 1)$	0.68	(0.19)
$\log(\text{number of claims})$	0.30	(0.03)
$\mathbf{1}(\text{revenue} = 0)$	1.42	(0.14)
$\log(\text{revenue})$	0.14	(0.02)
$\mathbf{1}(\text{employees} = 0)$	0.45	(0.07)
$\log(\text{employees})$	-0.01	(0.02)
Application year	-0.03	(0.05)
$(\text{Application year})^2$	-0.01	(0.01)
Decision year	0.30	(0.06)
$(\text{Decision year})^2$	-0.03	(0.01)
Constant	-1.40	(0.21)
$\log(\sigma)$	0.24	(0.05)
Observations	596	
Art units	260	
χ^2	10,353	

Notes: This table reports the relationship between KPSS ξ patent value, and patent application and firm level covariates. Coefficient estimates are from a Poisson model with art unit random effects. The sample is the subsample of granted patents for which the Kogan et al. (2017) measure of patent value is available in our analysis sample, except we retain firms with more than 100 million in 2014 revenue (unlike in our analysis sample) in order to maximize sample size (N=596). The dependent variable is the KPSS measure of patent value ξ in millions of 1982 dollars. Standard errors are reported in parentheses. Patent family size measures the number of countries in which the patent application was submitted. Number of claims measures the number of claims in the published US patent application. Revenue (in thousands of 2014 dollars) and number of employees are measured as of the year the US patent application was filed. $\log(\sigma)$ reports the log of the estimated variance of the art unit random effects. χ^2 reports the results of a likelihood ratio test statistic against a restricted Poisson model without art unit random effects; this test has one degree of freedom.

Table 5: Impacts on Firm Aggregates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Positive employ- ment	Log firm size	Revenue per worker	Value added per worker	EBITD per worker	Wage bill per worker	Surplus per worker	Labor comp per worker	W2 + 1099 per worker	Income tax per worker
Panel A: all firms										
High value (<i>Q5</i>)	0.00 (0.04)	0.22 (0.09)	36.75 (14.92)	15.74 (5.25)	9.11 (3.83)	3.65 (1.55)	12.41 (3.56)	3.94 (2.79)	2.80 (1.53)	0.77 (0.68)
Mean of outcome (<i>Q5</i>)	0.69	2.84	286.00	110.70	9.07	52.49	62.72	51.25	45.48	16.11
% Impact (<i>Q5</i>)	-0.6		12.8	14.2	100.4	7.0	19.8	7.7	6.2	4.8
Lower value (< <i>Q5</i>)	0.00 (0.01)	0.03 (0.04)	-9.68 (8.40)	0.84 (3.82)	-1.42 (1.77)	0.80 (0.90)	-0.26 (2.05)	1.32 (1.65)	0.52 (0.86)	0.74 (0.56)
Observations	155,646	103,437	103,437	103,437	103,437	103,437	103,437	103,437	107,789	103,159
Panel B: closely held firms										
High value (<i>Q5</i>)	-0.04 (0.08)	0.33 (0.14)	79.75 (24.60)	32.72 (8.24)	21.84 (6.08)	8.49 (3.42)	31.11 (7.44)	7.13 (1.38)	5.17 (1.98)	2.21 (1.64)
Mean of outcome (<i>Q5</i>)	0.70	2.80	283.80	111.80	18.67	46.47	66.17	44.22	40.53	16.34
% Impact (<i>Q5</i>)	-5.2		28.1	29.3	117.0	18.3	47.0	16.1	12.8	13.6
Lower value (< <i>Q5</i>)	-0.02 (0.02)	0.02 (0.05)	-3.29 (13.95)	-0.70 (6.19)	0.71 (3.19)	0.30 (1.51)	0.53 (4.03)	0.34 (2.55)	-0.11 (1.41)	1.23 (1.07)
Observations	75,132	49,943	49,943	49,943	49,943	49,943	49,943	49,943	51,998	49,808

Notes: This table reports difference-in-differences estimates of the effect of initial patent allowances on firm and worker outcomes, separately for high and low ex-ante valuable patent applications, in our analysis sample. Panel B replicates Panel A, but restricts the analysis to S-corporations and partnerships (“closely held” firms). Estimates correspond to coefficients on interactions of the designated value category with a post-decision indicator and an indicator for the application being initially allowed. Controls include main effect of value category interacted with a post-decision indicator, firm fixed effects, and art unit by application year fixed effects, as in equation (8). Standard errors (reported in parentheses) are two-way clustered by (1) art unit, and (2) application year by decision year. Panel B documents impacts on closely held firms. EBITD is earnings before interest, taxes, and deductions. Surplus is EBITD + wage bill. Labor compensation measures total deductions for labor expenses claimed by the firm. “W2 + 1099” measures to the sum of W2 and 1099 earnings divided by the sum of the number of W2’s and 1099’s filed. “Income tax per worker” is the average worker’s individual income tax liability. Revenue, value added, EBITD, wage bill, surplus, labor compensation, and W2 + 1099 pay are reported in thousands of 2014 USD.

Table 6: Workforce Composition

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Share female	Share inventors	Avg entrant earnings (yr bef ent)	Avg separator earnings (yr bef sep)	Avg stayer earnings (in app yr)	Avg age	Log quality	Log quality (expanded)
High value ($Q5$)	-0.01 (0.01)	-0.01 (0.01)	-0.84 (2.05)	0.72 (1.11)	1.29 (1.58)	-1.10 (0.56)	-0.02 (0.03)	-0.01 (0.03)
Mean of outcome ($Q5$)	0.31	0.09	24.42	28.19	66.59	41.74	10.43	10.56
% Impact ($Q5$)	-1.9	-12.5	-3.4	2.6	1.9	-2.6	-0.2	-0.1
Lower value ($< Q5$)	-0.01 (0.01)	-0.01 (0.01)	0.49 (0.70)	0.00 (0.53)	1.01 (1.19)	0.08 (0.22)	0.00 (0.01)	-0.01 (0.01)
Observations	103,437	103,437	70,079	75,524	99,558	103,434	103,437	97,786

Notes: This table reports difference-in-differences estimates of the effect of initial patent allowances on within-firm workforce composition measures, separately for high and low ex-ante valuable patent applications, in our analysis sample. Estimates correspond to coefficients on interactions of the designated value category with a post-decision indicator and an indicator for the application initially allowed. Controls include main effect of value category interacted with a post-decision indicator, firm fixed effects, and art unit by application year by calendar year fixed effects, as in equation (8). Standard errors (reported in parentheses) are two-way clustered by (1) art unit, and (2) application year by decision year. “Avg entrant earnings (yr bef ent)” measures the earnings of entrants in the year before they joined the firm. “Avg separator earnings (yr bef sep)” measures the earnings of separators in the year before they leave the firm. Stayers are defined as those who were employed by the same firm in the year of application. Entrants are defined as those employees who were not employed at the firm in the previous year. “Avg age” measures the average age of all employees at the firm. Earnings are measured in thousands of 2014 USD. “Log quality” gives predicted log wage based on worker demographics and inventor status. “Log quality (expanded)” includes wage on previous job in predictive model. See text for details.

Table 7: Earnings Impacts by Year of Entry/Exit

	Change since application year							
	(1) Avg cohort earnings	(2) Avg stayer earnings	(3) Avg leaver earnings	(4) Avg entrant earnings	(5) Avg recent entrant earnings	(6) Avg stayer earnings	(7) Avg leaver earnings	(8) Avg entrant earnings
High value (<i>Q5</i>)	3.96 (2.29)	7.78 (2.93)	-1.54 (1.94)	0.11 (1.64)	-2.71 (1.81)	6.50 (3.10)	2.77 (5.65)	0.95 (1.80)
Mean of outcome (<i>Q5</i>)	57.39	72.56	50.57	30.13	38.93	72.56	50.57	30.13
% Impact (<i>Q5</i>)	6.9	10.7	-3.0	0.4	-7.0	9.0	5.5	3.2
Lower value (< <i>Q5</i>)	0.34 (1.18)	2.48 (1.59)	0.90 (1.39)	0.31 (0.79)	0.78 (1.01)	1.48 (1.63)	-3.87 (2.40)	-0.18 (0.70)
Observations	151,892	99,558	109,169	70,079	68,691	99,558	109,169	70,079

Notes: This table reports difference-in-differences estimates of the effect of initial patent allowances on worker outcomes for employees who stay, enter, and exit $\hat{\xi}$, separately for high and low ex-ante valuable patent applications. Estimates correspond to coefficients on interactions of the designated value category with a post-decision indicator and an indicator for the application being initially allowed. Controls include main effect of value category interacted with a post-decision indicator, firm fixed effects, and art unit by application year by calendar year fixed effects, as in equation (8). Standard errors (reported in parentheses) are two-way clustered by (1) art unit, and (2) application year by decision year. “Avg cohort earnings” measures the W2 earnings of workers employed by the firm in the year of application. “Avg stayer earnings” measures the W2 earnings of workers employed by the firm in the year of application who are also employed in the present year. “Avg leaver earnings” measures the W2 earnings of workers employed by the firm in the year of application who are not employed in the present year. “Avg entrant earnings” measures the average earnings of employees hired by the firm within the past three years. “Change since application year” columns are earnings in the current year (Columns 6 and 7), and year of entry (Column 8) minus the respective values in the application year. Stayers are defined as those who were employed by the same firm in the year of application. Entrants are defined as those employees who were not employed at the firm in the previous year. Earnings are measured in thousands of 2014 USD.

Table 8: Pass-Through Estimates

	Wage bill per worker	Avg male earnings	Avg non-inventor earnings	Avg stayer earnings	Avg stayer earnings minus non-inventor stayer earnings	Avg non-inventor stayer earnings						
	(1a) OLS	(2a) OLS	(3a) OLS	(4a) OLS	(5a) OLS	(6a) OLS						
	(1b) IV	(2b) IV	(3b) IV	(4b) IV	(5b) IV	(6b) IV						
Panel A: Surplus / worker												
Impact	0.16 (0.01)	0.29 (0.12)	0.18 (0.01)	0.53 (0.18)	0.13 (0.00)	0.19 (0.01)	0.20 (0.01)	0.61 (0.30)	0.19 (0.01)	0.51 (0.27)	0.16 (0.00)	0.48 (0.22)
Elasticity	0.19	0.35	0.19	0.54	0.17	0.25	0.19	0.56	0.26	0.70	0.17	0.50
Observations	103,437	103,437	95,004	95,004	100,901	100,901	99,558	99,558	99,558	99,558	94,909	94,909
1 st stage F	12.12	10.60	9.34	13.38	9.34	13.38	13.38	13.38	13.38	13.38	13.38	8.93
Exogeneity	0.29	0.08	0.60	0.14	0.60	0.14	0.14	0.14	0.22	0.22	0.08	0.08
Anderson-Rubin 90% CI	(0.10,0.57)	(0.27,0.98)	(-0.01,0.43)	(0.21,1.36)	(-0.01,0.43)	(0.21,1.36)	(0.11,1.14)	(0.21,1.18)	(0.11,1.14)	(0.08,1.15)	(0.18,1.17)	(0.21,1.18)
Panel B: Value added / worker												
Impact	0.07 (0.00)	0.23 (0.13)	0.08 (0.00)	0.38 (0.19)	0.06 (0.00)	0.15 (0.10)	0.08 (0.00)	0.49 (0.28)	0.06 (0.00)	0.41 (0.25)	0.06 (0.00)	0.41 (0.20)
Elasticity	0.15	0.47	0.14	0.67	0.13	0.33	0.13	0.78	0.08	0.57	0.11	0.72
Observations	103,437	103,437	95,004	95,004	100,901	100,901	99,558	99,558	99,558	99,558	94,909	94,909
1 st stage F	8.99	8.99	8.27	8.27	8.79	8.79	9.39	9.39	9.39	9.39	6.86	6.86
Exogeneity	0.17	0.17	0.08	0.08	0.35	0.08	0.08	0.08	0.11	0.11	0.03	0.03
Anderson-Rubin 90% CI	(0.07,0.61)	(0.15,1.03)	(-0.01,0.38)	(0.15,1.39)	(-0.01,0.38)	(0.15,1.39)	(0.08,1.15)	(0.18,1.17)	(0.08,1.15)	(0.18,1.17)	(0.18,1.17)	(0.18,1.17)

Notes: This table reports OLS and IV estimates of the effect of increases in surplus per worker on selected earnings outcomes. The excluded instrument is the interaction of top quintile of ex-ante value ξ category with a post-decision indicator and an indicator for the application being initially allowed. Controls include main effect of value category interacted with post-decision indicator and interaction of lower quintile value category with a post-decision indicator interacted with an indicator for initially allowed, firm fixed effects, and art unit by application year by calendar year fixed effects. Standard errors (reported in parentheses) are two-way clustered by (1) art unit, and (2) application year by decision year. “Exogeneity” reports the p-value from a test of the null hypothesis that IV and OLS estimators have same probability limit. Stayers are defined as those who were employed by the same firm in the year of application. Surplus is EBITD (earnings before interest, tax, and depreciation) + wage bill. Earnings, wage bill, and surplus are measured in thousands of 2014 USD.

Table 9: Retention of application cohort

	(1)	(2)	(3)	(4)	(5)
	All	Above median	Men	Women	Non-inventors
Retention elasticity	1.22 (0.58)	1.41 (0.65)	0.80 (0.35)	1.17 (0.80)	1.31 (0.68)
Separation elasticity	-1.74	-2.91	-1.23	-1.88	-1.79
Observations	99,558	81,728	88,100	71,591	94,909
1 st stage F	7.81	5.80	31.13	3.61	6.74
Exogeneity	0.034	0.029	0.041	0.060	0.047
Anderson-Rubin 90% CI	(0.459, 3.080)	(0.597, 4.091)	(0.283, 1.524)	(0.233, 8.687)	(0.422, 3.655)

Notes: This table reports IV estimates of the effect of increases in selected earnings measures on the retention of employees. The excluded instrument is the interaction of top quintile of ex-ante value $\hat{\xi}$ category with a post-decision indicator and an indicator for the application being initially allowed. Controls include main effect of value category interacted with post-decision indicator and interaction of lower quintile value category with a post-decision indicator interacted with an indicator for initially allowed, firm fixed effects, and art unit by application year by calendar year fixed effects. Standard errors (reported in parentheses) are two-way clustered by (1) art unit, and (2) application year by decision year. “Separation elasticity” reports the elasticity of the outcome retention measure with respect to the associated earnings measure at the mean of the outcome and the mean earnings measure. “Exogeneity” reports p-value for test of null hypothesis that IV and OLS estimators have same probability limit. “Above median” refers to members of the application cohort who earned above that firm’s median in the application year. Stayers are defined as those who were employed by the same firm in the year of application. Earnings are measured in thousands of 2014 USD.

A Appendix: Data

A.1 Description of patent data

Our patent data build draws on several sources. Three identification numbers are relevant when using these datasets. First, publication numbers are unique identifiers assigned to published patent applications. Second, application numbers are unique identifiers assigned to patent applications that in practice are quite similar to publication numbers, but sometimes one application number is associated with multiple publication numbers. Finally, patent grant numbers are unique identifiers assigned to granted patents. Note that one patent application number can be associated with more than one granted patent.

Traditionally, unsuccessful patent applications were not published by the USPTO. However, as part of the American Inventors Protection Act of 1999, the vast majority of patent applications filed in the US on or after 29 November 2000 are published eighteen months after the filing date. There are two exceptions. First, applications granted or abandoned before eighteen months do not appear in this sample unless the applicant chooses to ask for early publication. Lemley and Sampat (2008) estimate that about 17 percent of patents are granted before eighteen months, of which about half (46 percent) are published pre-patent grant. Second, applications pending more than eighteen months can “opt out” of publication if they do not have corresponding foreign applications, or if they have corresponding foreign applications but also have priority dates predating the effective date of the law requiring publication (Lemley, and Sampat 2008).²⁹

1. Census of published USPTO patent applications. We observe the census of published (accepted and rejected) patent applications published by the US Patent and Trademark Office (USPTO). Our source for this data is a set of bulk XML files hosted by Google.³⁰ The underlying XML file formats were often inconsistent across years, so in the process of parsing these XML files to flat files we attempted to validate the data against other USPTO administrative data wherever possible. These records are at the publication number level.
2. Census of granted USPTO patents. For the published USPTO patent applications in our data, we wish to observe which of those applications were granted patents. Our source for this data is a set of bulk XML files hosted by Google.³¹ As with the published USPTO patent applications data, the underlying XML file formats were often inconsistent across years, so in the process of parsing these XML files to flat files we attempted to validate the data against other USPTO administrative data wherever possible. As one specific example, even though patent numbers uniquely identify patent grants, there are twenty-one patent numbers in this data that appear in the data twice with different grant dates. Checking these patent numbers on the USPTO’s online Patent Full Text (PatFT) database reveals that in each of these cases, the duplicated patent number with the earlier grant date is correct.³² Accordingly, we drop the twenty-one observations with the later grant dates.
3. USPTO patent assignment records. Some of our published patent applications are missing assignee information. (Applicants are not required to submit assignee information to the USPTO at the time of application.) Based on informal conversations with individuals at the USPTO, we fill in missing assignee names to the extent possible using the USPTO Patent Assignment data. The USPTO Patent Assignment data records assignment transactions, which are legal transfers of all or part of the right, title, and interest in a patent or application from one or more existing owner to one or more recipient. The dataset is hosted on the USPTO

²⁹For more details, see <http://www.uspto.gov/web/offices/pac/mpep/s1120.html> and the discussion in Lemley and Sampat (2010). Most applications not published eighteen months after filing are instead published sixty months after filing.

³⁰See <http://www.google.com/googlebooks/uspto-patents-applications-biblio.html>.

³¹See <http://www.google.com/googlebooks/uspto-patents-grants-text.html>.

³²PatFT can be accessed at <http://patft.uspto.gov/>.

website.³³ Each transaction is associated with a patent number, application number, and/or publication number (wherever each is applicable). The patent assignment records include both initial assignments and re-assignments, but only the former is conceptually appropriate for our analysis since we want to measure invention ownership at the time of application. We isolate initial assignments by taking the assignment from this database with the earliest execution date. If a given assignment has more than one execution date (e.g., if the patent application is assigned to more than one entity), we use the latest execution date within that assignment as the transaction execution date. Using these initial assignments, we fill in assignee organization name as well as assignee address information where possible when these variables are missing from our published patent applications data.

4. USPTO patent document pre-grant authority files. A very small number (1,025 total) of published USPTO patent applications are “withdrawn,” and these observations tend to be inconsistently reported across the various datasets we analyze. The USPTO patent document pre-grant authority files — an administrative data file hosted on the USPTO website — allows us to exclude all withdrawn applications for consistency.³⁴ Our versions of these files were downloaded on 24 March 2014 and are up to date as of February 2014. These records are at the publication number level.
5. USPTO PAIR records. We analyze several variables, such as the date of initial decisions, from the USPTO Patent Application Information Retrieval (PAIR) data, which we draw from an administrative dataset called the Patent Examination Research Dataset (PatEx).³⁵ With the exception of 264 published patent applications, these data are available for our full sample of published USPTO patent applications. These records are at the application number level.
6. Examiner art unit and pay scale data. Frakes and Wasserman (2017) generously provided us with examiner art unit and General Schedule (GS) pay scale data they received through FOIA requests. These data allow us to identify which examiners were active in each art unit in each year.
7. Thomson Innovation database. All of the databases listed above record information obtained directly from the USPTO. One measure of patent value that cannot be constructed based on the USPTO records alone is a measure of patent family size, as developed in Jonathan Putnam’s dissertation (Putnam 1996). Generally stated, a patent “family” is defined as a set of patent applications filed with different patenting authorities (e.g., the US, Europe, Japan) that refer to the same invention. The key idea is that if there is a per-country cost of filing for a patent, firms will be more likely to file a patent application in multiple countries if they perceive the patent to have higher private value. Past work — starting with Putnam (1996) — has documented evidence that patent family size is correlated with other measures of patent value. The Thomson Reuters Innovation database collects non-US patent records, and hence allows for the construction of such a family size measure.³⁶ We purchased a subscription to the Thomson Innovation database, and exported data from the web interface on all available variables for all published USPTO patent applications. To construct our family size measure, we take the DWPI family variable available in the Thomson Innovation database (which lists family members), separate the country code from the beginning of each number (e.g., “US” in “US20010003111”), and then count the number of unique country codes in the family. These records are at the publication number level.
8. Hall, Jaffe, and Trajtenberg (2001) NBER data. Hall, Jaffe, and Trajtenberg (2001) constructed a match between US patents granted between January 1963 and December 1999 with the Compustat data. As part

³³Available at: <https://www.uspto.gov/learning-and-resources/electronic-data-products/patent-assignment-dataset>.

³⁴See <http://www.uspto.gov/patents/process/search/authority/>.

³⁵See <http://www.uspto.gov/learning-and-resources/electronic-data-products/patent-examination-research-dataset-public> for the underlying PAIR data, see: <http://portal.uspto.gov/pair/PublicPair>.

³⁶See <http://info.thomsoninnovation.com/>.

of that work, the authors constructed technology categories to describe the broad content area of different patents, based on categorizations of the patent technology class and subclass variables.³⁷ We match on these technology categories, and hand-fill the small number of cases in which classes or subclasses appear in our data but not in the crosswalk constructed by Hall and co-authors. These records are at the patent class level.

9. Kogan et al. (2017) patent value data. Kogan et al. (2017) provide their final estimates of patent value for their sample of granted patents at <https://iu.app.box.com/v/patents/>. In particular we downloaded the “patents.zip” file, which contains a linkage between USPTO patent grant numbers and the estimate of the patent value ξ . These data were downloaded on — and are accurate as of — 7 August 2016. To develop a measure of patent value at the application number level, we associate each application with its potentially numerous patent numbers. We then sum the values of ξ by application number to obtain a measure of the ex-post value of granted applications.
10. USPTO technology center data. Technology centers are groupings of examiner art units. The USPTO hosts a listing of all technology centers and associated examiner art units at <http://www.uspto.gov/patent/contact-patents/patent-technology-centers-management>. We use these groupings to examine heterogeneity in predicted patent value by area of invention in Appendix Table C.2.

A.2 Construction of patent application sample

We restrict the sample to USPTO patent applications filed on or after 29 November 2000 (the date when “rejected” applications started to be published), and ends with applications published on 31 December 2013. We impose a few additional sample restrictions:

- We exclude a very small number of “withdrawn” patent applications (1,025 total) given that these observations tend to be inconsistently reported across datasets. As noted above, the withdrawn applications were identified using the USPTO patent document pre-grant authority files.
- Six publication numbers are listed in the USPTO patent document pre-grant authority files but are not available in any of our other datasets;³⁸ we exclude these observations from our sample.
- Four publication numbers are missing from the Thomson Innovation database.³⁹ We include these observations in the sample, but they are missing data for all variables drawn from the Thomson Innovation data.
- Based on the kind code variable listed on the USPTO published patent applications,⁴⁰ we exclude a small number of patent applications that are corrections of previously published applications: corrections of published utility/plant patent applications (kind codes A9/P9; 3,156 total), and second or subsequent publications of the same patent application (kind codes A2/P4; 1,182 total). These kind codes more generally allow us to confirm that our sample does not include various types of documents: statutory invention registration documents (kind code H1), reexamination certificates (kind codes Bn/Cn/Fn for n=1-9), post grant review certificates (kind codes Jn for n=1-9), inter parties review certificates (kind codes Kn for n=1-9), or derivation certificates (kind codes On for n=1-9). Our final sample includes only utility patent applications (kind code A1; 3,597,787 total) and plant patent applications (kind code P1; 4,196 total).

Finally, there are two data inconsistencies that we have resolved as follows:

³⁷See <http://www.nber.org/patents/>.

³⁸Specifically, these publication numbers are: US20010003111; US20020011585; US20020054271; US20020084413; US20020084764; and US20020103782.

³⁹Specifically, these publication numbers are: US20010020331; US20010020666; US20010021099; and US20010021102.

⁴⁰For a summary of USPTO kind codes, see: <http://www.uspto.gov/patents/process/search/authority/kindcode.jsp>.

- Seven observations appear to be missing from Google’s XML files of the published patent applications.⁴¹ We were able to hand-code the required variables for these observations based on the published patent applications posted at <http://patft.uspto.gov> for all but three of these observations (specifically, publication numbers US20020020603; US20020022313; US20020085735). For those three observations, we hand-coded the required variables based on the information available at <http://portal.uspto.gov/pair/PublicPair/>; for these, we assumed that the appropriate correspondent addresses were those listed in the “Address and Attorney/Agent” field under Correspondence Address.
- The applications data contain 67 applications that were approved SIR (statutory invention registration) status but have the kind code “A1,” instead of “H1” (as we would expect). We changed the kind code to “H1” for these applications, and they are therefore dropped from our sample.

A.3 Description of US Treasury tax files

All firm-level variables are constructed from annual business tax returns over the years 1997-2014: C-Corporations (Form 1120), S-Corporations (Form 1120S), and Partnerships (Form 1065). Worker-level variables are constructed from annual tax returns over the years 1999–2014: Employees (form W2) and contractors (form 1099).⁴²

Variable Definitions

To define firm-level variables using the US Treasury files, we use the following line items from the 2010 business tax forms: 1120 for C-corporations, 1120S for S-corporations, and 1065 for partnerships. Note that the tax form line numbers can sometimes change slightly if, for example, a line is added for a new deduction.

- Revenue
 - Line 1c of Form 1120 for C-Corporations, Form 1120S for S-Corporations, and Form 1065 for partnerships. When 1c is not available, we use 1a, which is gross receipts. We replace negative revenue entries, which are very rare, with missing values.
- Total Income
 - For C-Corporations, line 11 on Form 1120. Note that this subtracts COGS from revenues and includes income from a variety of sources (e.g., dividends, royalties, capital gains, etc). For S-Corporations, line 6 on Form 1120S. For partnerships, line 8 on Form 1065.
- Total Deductions
 - For C-corporations, line 27 on Form 1120. For S-corporations, line 20 on Form 1120S. For partnerships, line 21 on Form 1065.
- Labor Compensation
 - For C-Corporations, sum of lines 12, 13, 24, and 25 on Form 1120.⁴³ For S-Corporations, sum of lines 7, 8, 17, and 18 for Form 1120S. For partnerships, sum of lines 9, 10, 18, and 19 on Form 1065. These lines are compensation to officers, salaries and wages, retirement plans, and employment benefit programs, respectively.⁴⁴

⁴¹Specifically, the missing publication numbers are: US20010020331; US20010020666; US20010021099; US20010021102; US20020020603; US20020022313; US20020085735.

⁴²W2 data are not available in 1997–1999.

⁴³Ideally, we could also add Schedule A line 3, which is the cost of labor on the COGS Form 1125-A, but these data are not available. However, the W2-based measure of compensation avoids this issue.

⁴⁴For partnerships, the compensation to officers term is called “Guaranteed payments to partners.”

- Value Added
 - Gross receipts minus the difference between cost of goods sold and cost of labor.
 - For C-Corporations, line 3 on Form 1120. For S-corporations, line 3 on Form 1120S. For partnerships, line 3 on Form 1065.⁴⁵
- Profits
 - Yagan (2015) defines operating profits as revenues less Costs of Goods Sold and deductions where deductions are total deductions other than compensation to officers, interest expenses, depreciation, and domestic production activities deduction. We do not add back compensation to officers.
 - For C-Corporations, we define operating profits as the sum of lines 1c, 18, and 20, less the sum of 2 and 27 on Form 1120. We set profits to missing if 1c, 18, 20, 2, and 27 are all equal to zero.
 - For S-Corporations, operating profits are the sum of lines 1c, 13, and 14 less the sum of 2 and 20 on Form 1120S.
 - For partnerships, operating profits are the sum of lines 1c, 15, and 16c less the sum of 2 and 21 on Form 1065.
- EBITD
 - EBITD is total income less total deductions (other than interest and depreciation).
 - For C-Corporations, it is the sum of lines 11, 18, and 20, less 27 on Form 1120.
 - For S-Corporations, it is the sum of lines 1c, 13, and 14 less 20 on Form 1120S.
 - For partnerships, it is the sum of lines 1c, 15, and 16c less 21 on Form 1065.
- Employment
 - Number of W2s associated with an Employer Identification Number (EIN).
- Wage bill per worker
 - Sum of W2 box 1 payments divided by number of W2s for a given EIN.
- Surplus
 - Sum of EBITD and Wage bill, which is the sum of W2 box 1 payments for a given EIN.
- Inventor earnings per inventor
 - Wage bill per worker for workers who are identified as inventors by Bell et al. (2016).
- Cohort earnings per worker
 - Wage bill per worker for workers who were employed at the firm in the year of application regardless of whether or not they stay at the firm.
- Stayer earnings per worker

⁴⁵Line 3 is calculated as line 1c minus line 2.

- Stayers are cohort earning per worker for the set of workers who are still at the firm.
- Leaver earnings per worker
 - Leavers are cohort earning per worker for the set of workers in the initial cohort who are no longer at the firm, i.e., are no longer receiving a W2 associated with the original firm that applied for a patent.
- Earnings Gap Q4-Q1
 - Average earnings within quartile four and quartile one of a firm's wage distribution.
- Separators
 - The number of workers who left the EIN in the previous year.
- Entrants
 - The number of workers who joined the EIN relative to the previous year.
- State
 - Uses the state from the business's filing address.
- Entity Type
 - Indicator based on tax-form filing type.
- Industry
 - NAICS codes are line 21 on Schedule K of Form 1120 for C-Corporations, line 2a Schedule B of Form 1120S for S-Corporations, and Box A of Form 1065 for partnerships.
- Active Firm
 - An active firm has non-zero and non-missing total income and non-missing total deductions.

A.3.1 Deflator to convert to 2014 USD

Table A.1: Deflator to convert to 2014 USD

Year	1 / 2014 CPI	Year	1 / 2014 CPI
1993	1.503062988	2004	1.219663817
1994	1.47180133	2005	1.181649182
1995	1.441698831	2006	1.146416065
1996	1.415894851	2007	1.116642696
1997	1.391942424	2008	1.095506864
1998	1.377006398	2009	1.08694
1999	1.357571973	2010	1.073775512
2000	1.327317133	2011	1.052064076
2001	1.297761328	2012	1.033016537
2002	1.278151458	2013	1.016449245
2003	1.253173459	2014	1

Notes: This table shows the deflators used to convert our dollar amounts from current dollars into 2014 USD. Deflators were calculated using price data from the US Bureau of Economic Analysis (BEA), National Income and Product Accounts (NIPA) Table 1.1.4: ‘Price Indexes for Gross Domestic Product.’ See US Bureau of Economic Analysis (2014).

A.4 Description of merge between patent applications data and US Treasury tax files

Our analysis relies on a new merge between published patent applications submitted to the US Patent and Trademark Office (USPTO) and US Treasury tax files. Below we describe the details of this merge, which relies on a fuzzy matching algorithm to link USPTO assignee names with US Treasury firm names.

A.4.1 Creating standardized names within the patent data

Published patent applications list an assignee name, which reflects ownership of the patent application. Due to, e.g., spelling differences, multiple assignee names in the USPTO published patent applications data can correspond to a single firm. For example, “ALCATEL-LUCENT U.S.A., INC.,” “ALCATEL-LUCENT USA, INCORPORATED,” and “ALCATEL-LUCENT USA INC” are all assigned the standardized name “alcatel lucent usa corp”.

We employed a name standardization routine as follows. Starting with names in unicode format, we transform the text into Roman alphabet analogs using the “unidecode” library to map any foreign characters into their applicable English phonemes, and then shift all characters to lowercase.⁴⁶ We then standardize common terms that take multiple forms, such as “corp.” and “corporation”; these recodings were built on the name standardization routine used by the National Bureau of Economic Research (NBER)’s Patent Data Project, with modifications as we saw opportunities to improve that routine.⁴⁷ We additionally eliminate any English articles (such as “a” or “an”), since these appeared to be uninformative in our attempts to uniquely identify entities. We then tokenize standardized names by splitting on natural delimiters (e.g., spaces and commas), after which we remove any non-alphanumeric

⁴⁶The unidecode library is available at <https://github.com/iki/unidecode>, and is a direct Python port of the Text::Unidecode Perl module by Sean M. Burke.

⁴⁷The NBER Patent Data Project standardization code is available at <https://sites.google.com/site/patentdataproyect/Home/posts/namestandardizationroutinesuploaded>.

punctuation. Finally, sequences of single-character tokens are merged into a combined token (e.g., “3 m corp” would become “3m corp”). The resultant ordered list of tokens constitutes our standardized entity name. We refer to the USPTO standardized firm name as $SNAME_{USPTO}$.

A.4.2 Creating standardized names within the US Treasury tax files

In the US Treasury tax files, firms are indexed by their Employer Identification Number (EIN). Each EIN is required to file a tax return for each year that it is in operation. Specifically, we restrict our analysis to firms with valid 1120, 1120S, or 1065 filings over the years 1997–2014. We apply the same name standardization algorithm to the Treasury firm names that was applied to the USPTO names. We refer to the Treasury standardized firm name as $SNAME_{Treasury}$.

A.4.3 Merging standardized names across the USPTO data and the US Treasury tax files

We then conduct a fuzzy merge of $SNAME_{USPTO}$ to $SNAME_{Treasury}$ using the SoftTFIDF algorithm, which is described below. We use this algorithm to allocate each $SNAME_{USPTO}$ to a single $SNAME_{Treasury}$, provided that match quality lies within a specified tolerance. To choose the tolerance we used a hand coded match of applications to Compustat firms as a validation dataset (see Section A.4.4). The tolerance (and other tuning parameters) were chosen to minimize the sum of Type I and II error rates associated with matches to Compustat firms. The resulting firm-level dataset has one observation per $SNAME_{Treasury}$ in each year. However, there are some cases in which multiple EINs are associated with a given $SNAME_{Treasury}$. In those cases, we chose the EIN with the largest total income in the year of application in order to select the most economically active entity associated with that standardized name.⁴⁸

SoftTFIDF algorithm Our firm name matching procedure of name $a \in SNAME_{USPTO}$ to name $b \in SNAME_{Treasury}$ works as follows. Among all the words in all the firm names in $SNAME_{USPTO}$ that are close to a given word in b , we pick the word with the highest SoftTFIDF index value, which is a word-frequency weighted measure of similarity among words. We do this for each word in the firm’s name. For instance, American Airlines Inc would have three words. We then take a weighted-sum of the index value for each word in the firm name where the weights are smaller for frequent words like "Inc." This weighted sum is the SoftTFIDF value at the level of firm-names (as opposed to words in firm names). We assign a to the firm name b with the highest SoftTFIDF value above a threshold; otherwise, the name a is unmatched. Because of computational limitations, we limit comparisons to cases in which both a and b start with the same letter. Therefore we will miss any matches that do not share the same first letter. This subsection provides details on this procedure and example matches.

SoftTFIDF of firm names A score between groups of words X, Y is given by

$$\text{SoftTFIDF}(X, Y) := \sum_{w \in X} \text{weight}(w, X) \cdot \alpha(w, Y)$$

where $\text{weight}(w, X)$ is a word frequency-based importance weight and $\alpha(w, Y)$ is a word match score that uses a word similarity index. Specifically, the importance weight for the word w in the set of words Z is: $\text{weight}(w, Z) := \frac{\text{tfidf}(w, Z)}{\sqrt{\sum_{w' \in Z} \text{tfidf}(w', Z)^2}}$, where

- $\text{tfidf}(w, Z) := \text{tf}(w, Z) \times \text{idf}(w, \mathcal{L})$,
- $\text{tf}(w, Z) := \frac{n(w, Z)}{\sum_{w' \in Z} n(w', Z)}$,

⁴⁸For example, if two EINs shared the same standardized name $SNAME_{Treasury}$ but one EIN made 50 million in total income and the other showed three million in total income, we chose the EIN that earns 50 million.

- $\text{idf}(w, \mathcal{Z}) := \log \left(\frac{|\mathcal{Z}|}{|\{Z \in \mathcal{Z} \mid w \in Z\}|} \right)$,
- $n(w, Z)$ is the number of occurrences of word w in a set of words Z ,
- \mathcal{Z} is the set of all words in either $SNAME_{USPTO}$ or $SNAME_{IRS}$.

We compute the word match score $\alpha(w, Y)$ for words that are close to those in $SNAME_{USPTO}$. To determine which names are close, we use a Jaro-Winkler distance metric to measure the distance between two strings.

Jaro-Winkler metric of distance between strings We use this metric since it has been shown to perform better at name-matching tasks (Cohen, Ravikumar, and Fienburg 2003) than other metrics such as Levenshtein distance, which assigns unit cost to every edit operation (insertion, deletion, or substitution). A key component of the Jaro-Winkler metric is the Jaro metric. The Jaro metric depends on the length of $SNAME_{USPTO}$, the length of $SNAME_{Treasury}$, the number of shared letters, and the number of needed transpositions of shared letters.

Specifically, consider strings $s = s_1 \dots s_K$ and $t = t_1 \dots t_L$ and define $H = \frac{\min\{|s|, |t|\}}{2}$, which is half the smaller of K and L . We say a character s_i is **in common with** t if $\exists j \in [i - H, i + H]$ s.t. $s_i = t_j$. Let s', t' be the ordered sets of **in-common** characters (hence we will re-index). Then define $T_{s', t'} := \frac{1}{2} |\{i \mid s'_i \neq t'_i\}|$. The similarity metric is given by

$$\text{Jaro}(s, t) := \frac{1}{3} \cdot \left(\frac{|s'|}{|s|} + \frac{|t'|}{|t|} + \frac{|s'| - T_{s', t'}}{|s'|} \right).$$

The Jaro-Winkler metric is given by

$$\text{Jaro-Winkler}(s, t) := \text{Jaro}(s, t) + \frac{P'}{10} \cdot (1 - \text{Jaro}(s, t)),$$

where P as the longest common prefix of t and s and then $P' = \max\{P, 4\}$, which is the normalization used in Cohen, Ravikumar, and Fienburg (2003).

Word match score $\alpha(w, Z)$ We define the word match score as follows:

$$\alpha(w, Z) = \begin{cases} 0 & \text{if } \text{closest}(\theta, w, Z) = \emptyset \\ \max_{w' \in \text{closest}(\theta, w, Z)} \text{weight}(w', Z) \cdot \text{Jaro-Winkler}(w, w') & \text{otherwise} \end{cases}$$

where

$$\text{closest}(\theta, w, Z) := \{v \in Z \mid \forall v' \in Z, (\text{Jaro-Winkler}(w, v) \geq \text{Jaro-Winkler}(w, v')) \wedge \text{Jaro-Winkler}(w, v) > \theta\}.$$

In words, we select the word w that is the closest importance-weighted match among words that are close to the word w in Z given closeness threshold θ . The accuracy of this matching procedure, which has also recently been used by Feigenbaum (2016), will likely become clearer after reviewing the following examples and discussing how we selected the tuning parameters (such as the closeness threshold θ).

Example

USPTO Assignee Name	Compustat Firm Name (best match)	Match Score
angiotech pharmaceuticals corp	angiotech pharmaceuticals	.9982
assg brooks justin	brooks resources corp	.5857
hewlett packard development corp	hewlett packard corp	.8482
huawei device corp	huatue electronics corp	.0013
matsushita electric works corp	matson corp	.0012
olympus corp	olympus capital corp	.9109
safety crafted solutions corp	safety first corp	.3862
sc johnson home storage corp	sc holdings corp	.5144

This table provides a small sample of candidate matches from a USPTO to Compustat match.⁴⁹

A.4.4 Validation: Compustat-USPTO match

This section describes the hand matching process we used to determine the true mapping of USPTO names to Compustat names for a random sample of USPTO names. We describe the hand coding task and how we use the hand coded linkages to select the tuning parameters.

Hand coding tasks We hired several workers on Upwork (formerly Odesk) as well as University of Chicago undergraduates to hand match two lists of names. The goal for these workers was to match every name in a source file (a list of 100 randomly selected USPTO names) to a target file of Compustat names or to conclude that there is no matching name in the target file. To increase accuracy, we informed these workers that (1) we had hand-coded several of these names ourselves, (2) every name in the source file would be assigned to multiple workers, and (3) we would only accept reasonably accurate work. We also instructed them to use Google to confirm that matches were true matches. For example, “infinity bio ltd” may seem like a match with “infinity pharmaceuticals inc,” but Googling the first reveals that the former is a small Brazilian energy company while the latter is a pharmaceutical company headquartered in the US. If one worker found a match but another did not, we considered the non-empty match to be correct. Overall, we ended up assigning 2,196 assignee names to workers, of which 286 (13%) had matches in the Compustat data.⁵⁰

Using hand coding tasks to select tuning parameters We use these hand-coded linkages to establish the “true mapping” from USPTO names to Compustat names, which enables us to select tuning parameters that minimize the sum of type I and type II errors (relative to these “true linkages”).

We constructed a grid and for each set of parameters on the grid executed a match. We then compared these fuzzy matchings to the “true mapping.” Type I errors occur when SoftTFIDF returns a match but either (a) the match is inconsistent with the hand-coded match or (b) the hand coded linkage shows no match at all. Type II errors occur when SoftTFIDF does not return a match but the hand-code process had a match.

The parameters that minimize the sum of these false positive and false negatives are: $\theta = .95$, token type of standardized names (instead of raw names), $P = 0$, and a threshold match score of .91. We remind the reader that the parameter θ governs the threshold similarity for two words to be considered “close.” Only “close” words contribute to a match score, hence $\theta = .95$ sets a relatively high cutoff below which two similar words do not increase the match score between two firm names. The prefix $P = 0$ suggests that not boosting scores by a common prefix doesn’t improve performance, which makes sense given that we block by the first letter already.⁵¹ Finally,

⁴⁹We present results using Compustat names instead of Treasury names for disclosure reasons.

⁵⁰This match rate is sensible: the number of Compustat names is roughly 20% the number of assignee names, so this rate is consistent with a reasonable proportion of Compustat firms applying for a patent.

⁵¹It is computationally infeasible to compare every single entity against every other, so we utilized first-letter blocking in order to reduce the sizes of sets being compared against one another. In particular, the target names (either Compustat or Treasury names) are chunked by

the threshold match score of .91 shows that we should only consider names a match if they are very close by our similarity metric. With these parameters, Type I and II errors are each below six percent.

A.4.5 Validation: Individual-inventor match

The Bell et al. (2016) inventor-level merge between patent applications and W2 reports in theory can — via the EINs provided on W2 reports — provide a linkage between patent applications and firms, but ex-ante we expect this inventor-based match to measure something conceptually different from a firm-based match. For example, many inventors work at firms that are not the assignee of their patents, in which case we would not expect our assignee-based merge to match to the same EIN as the Bell et al. (2016) inventor-based merge. However, the Bell et al. (2016) merge nonetheless provides a very valuable benchmark for assessing the quality of our assignee-based merge. Bell et al. generously agreed to share their inventor-based merge with us, and our preliminary results comparing the two linkages provide a second set of evidence supporting the quality of our assignee-based linkage. In the simplest comparison, around 70% of patent applications are associated with the same EIN in the two linkages. The characteristics of this match also look sensible, e.g., the match rates are higher if we limit the sample to patent applications that Bell et al. (2016) match to inventors who all work at the same firm. Given that we do not expect a match rate of 100% for the reasons detailed above, we view the results of this second validation exercise as quite promising.

the first letter of their standardized names and grouped with the source (USPTO) names with the same first letter. Hence, we will miss any matches that differ on the first letter.

B Appendix: Poisson model of patent value

Recall that the probability mass function for a Poisson distributed outcome Y with mean λ can be written:

$$p(Y|\lambda) = \exp(Y\lambda - \exp(\lambda)) / Y!$$

Let $Y_a = (Y_1, \dots, Y_{m_a})$ and $X_a = (X_1, \dots, X_{m_a})$ denote the vectors of outcomes and covariates respectively in an art unit a . Supposing $Y_j|X_a \stackrel{iid}{\sim} \text{Poisson}(X'_j\delta + v_a)$ where v_a is a scalar art unit effect, we can write:

$$\begin{aligned} \ln p(Y_a|X_a, v_a) &= \sum_{j=1}^{m_a} \ln p(Y_j|X'_j\delta + v_a) \\ &= \sum_{j=1}^{m_a} Y_j (X'_j\delta + v_a) - \exp(X'_j\delta + v_a) - \ln(Y_j!) \end{aligned}$$

The random effects Poisson likelihood of an art unit a can be written:

$$L(Y_a|X_a) = \frac{1}{\sqrt{2\pi\sigma_\eta}} \int \exp \left\{ \ln p(Y_a|X_a, v) - \frac{1}{2} \frac{v^2}{\sigma_v^2} \right\} dv$$

By independence across art units, the full log likelihood can be written $\sum_a \ln L(Y_a|X_a)$.

The first order condition for the coefficient vector δ is:

$$\begin{aligned} \frac{d}{d\delta} \sum_a \ln L(Y_a|X_a) &= \sum_a \frac{\int \left(\sum_{j=1}^{m_a} [Y_j - \exp(X'_j\delta + v)] X_j \right) \exp \left\{ \ln p(Y_a|X_a, v) - \frac{1}{2} \frac{v^2}{\sigma_v^2} \right\} dv}{\int \exp \left\{ \ln p(Y_a|X_a, v) - \frac{1}{2} \frac{v^2}{\sigma_v^2} \right\} dv} \\ &= \sum_a \sum_{j=1}^{m_a} \left[Y_j - \int \exp(X'_j\delta + v) \omega_a(v) dv \right] X_j = 0 \end{aligned}$$

where the weighting function $\omega_a(z) = \frac{\exp \left\{ \ln p(Y_a|X_a, z) - \frac{1}{2} \frac{z^2}{\sigma_v^2} \right\}}{\int \exp \left\{ \ln p(Y_a|X_a, v) - \frac{1}{2} \frac{v^2}{\sigma_v^2} \right\} dv}$ is the posterior density of v given the observables in art unit a . Note that this is a shrunken version of the usual Poisson orthogonality condition that is robust to misspecification of features of the conditional distribution other than the mean (Wooldridge 2010). The weights, however, rely on the exponential nature of the Poisson density function which, if misspecified, will yield inconsistency in small art units. In large art units, however, the posterior will spike around the “fixed effect” estimate of v , which is again robust to misspecification of higher moments of the conditional distribution.


The first order condition for the variance σ_v is:

$$\begin{aligned} \sum_a \frac{d}{d\sigma_v} \ln L(Y_a|X_a) &= \sum_a \left\{ -\frac{1}{\sigma_\eta} + \frac{\int \frac{\eta^2}{\sigma_\eta^3} \exp \left\{ \ln p(Y_a|X_a, v) - \frac{1}{2} \frac{v^2}{\sigma_v^2} \right\} dv}{\int \exp \left\{ \ln p(Y_a|X_a, v) - \frac{1}{2} \frac{v^2}{\sigma_v^2} \right\} dv} \right\} \\ &= \frac{1}{\sigma_\eta^3} \sum_a \left[\int \omega_a(v) v^2 dv - \sigma_v^2 \right] = 0. \end{aligned}$$

This simply says that the posterior variance of v in each art unit should average across art units to σ_v^2 .


C Appendix: Additional figures and tables

Figure C.1: Years Until Initial Decision




Notes: This figure plots a histogram of the years until the initial patent application decision for the sample of patent assignees by application pairs in the bottom row of Panel A of Table 1 (N=99,871).

Figure C.2: Years Until Patent Grant for Initially Rejected Patent Applications



Notes: This figure plots a histogram of the years until a patent grant for the subsample of patent assignees by application pairs in the bottom row of Panel A of Table 1 (N=99,871) which receive an initial rejection (N=88,298).

Figure C.3: Industry Composition of Firms



Notes: This figure plots the distribution of firms in our sample by industry. “Winners” are firms whose patent application is initially granted.

Table C.1: Testing for Spatial Correlation in Initial Allowance Decisions

	Initially allowed							
	State		Zip		4-D NAICS		4-D NAICS \times State	
	(1a)	(1b)	(2a)	(2b)	(3a)	(3b)	(4a)	(4b)
Panel A: analysis sample								
Intra-class correlation (ρ)	0.000	0.000	0.068	0.058	0.000	0.000	0.000	0.000
p-value	1.000	1.000	0.297	0.465	1.000	1.000	1.000	1.000
Observations	9,732	8,647	9,732	8,647	9,732	8,647	9,732	8,647
Number of categories	51	51	4,501	4,231	355	347	3,376	3,185
AU-AY FEs		✓		✓		✓		✓
Panel B: top quintile								
Intra-class correlation (ρ)	0.113	0.132	0.000	0.000	0.000	0.000	0.000	0.046
p-value	0.165	0.184	1.000	1.000	1.000	1.000	1.000	0.416
Observations	1,946	1,666	1,946	1,666	1,946	1,666	1,946	1,666
Number of categories	49	49	1,252	1,113	250	242	1,119	1,013
AU-AY FEs		✓		✓		✓		✓

Notes: This table reports the results of tests for whether initial patent allowances are geographically clustered, separately for the analysis sample and for the top-quintile predicted patent value sample. “Intra-class correlation (ρ)” reports the ratio of a random effects estimate of the geographic variance component to the sum of the geographic and idiosyncratic variance components. The p-value reports a Breusch-Pagan Lagrange multiplier test of the null hypothesis that $\rho=0$. AU-AY FEs denotes the inclusion of Art Unit (AU) by application year (AY) fixed effects.

Table C.2: Mean $\hat{\xi}$ by Technology Center

Technology center	$\bar{\hat{\xi}}$	N	Technology center	$\bar{\hat{\xi}}$	N
Business Methods - Finance	15.079	152	Telecomms: Analog Radio	3.080	43
Electronic Commerce	10.237	365	Mining, Roads, & Petroleum	2.991	518
Databases & File Mgmt	9.726	261	Microbiology	2.983	83
Tires, Adhesives, Glass, & Plastics	8.035	134	Semiconductors, Circuits, & Optics	2.903	237
2180: Computer Architecture	8.029	68	Molec Bio & Bioinformatics	2.891	68
Combust & Fluid Power Systems	7.803	111	Amusement & Education Devices	2.780	236
Aero, Agriculture, & Weaponry	7.129	224	Static Structures & Furniture	2.627	560
Selective Visual Display Systems	6.387	200	Fuel Cells & Batteries	2.437	177
Computer Graphic Processing	6.012	299	Business Methods	2.416	193
Optics	5.650	341	2110: Computer Architecture	2.163	50
Organic Chemistry & Polymers	5.622	116	Software Development	2.141	76
Organic Compounds	5.316	115	Medical Instruments	2.110	132
Organic Chemistry	4.913	62	Multiplex & VoIP	1.990	72
Manufact Devices & Processes	4.870	443	Metallurgy & Inorganic Chemistry	1.980	102
Memory Access & Control	4.865	49	Chemical Apparatus	1.947	171
Selective Communication	4.674	294	Semiconductors & Memory	1.936	185
Surface Transportation	4.428	294	Cryptography & Security	1.898	76
Electrical Circuits & Systems	4.090	266	Medical & Surgical Instruments	1.823	118
Coating, Etching, & Cleaning	3.615	83	Computer Networks	1.733	145
Misc. Computer Applications	3.459	129	Radio, Robotics, & Nucl Systems	1.597	85
Material & Article Handling	3.248	255	Receptacles, Shoes, & Apparel	1.444	470
Graphical User Interface	3.217	152	Kinestherapy & Exercising	1.330	138
Refrigeration & Combustion	3.084	265	Fluid Handling	0.706	188

Notes: This table reports the mean predictions of ex-ante value $\hat{\xi}$ by USPTO technology center of the application; technology centers are administrative groupings of art units designated by the USPTO. The sample is observations from our analysis sample whose application belongs to a technology center with more than 20 observations in the analysis sample (N=6,402). $\hat{\xi}$ is measured in millions of 1982 USD.

Table C.3: Heterogeneity across larger and smaller firms

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Positive employ- ment	log(Firm size)	Revenue per worker	Value added per worker	EBITD per worker	Wage bill per worker	Surplus per worker	Labor comp per worker	W2 + 1099 per worker	Average stayer earnings
Panel A: large firms										
High value (<i>Q5</i>)	0.05 (0.03)	0.20 (0.07)	39.36 (17.51)	14.04 (6.95)	8.89 (5.38)	2.16 (2.73)	10.67 (6.26)	1.69 (3.74)	1.98 (2.10)	6.81 (3.13)
Mean of outcome (<i>Q5</i>)	0.74	3.87	263.70	96.66	7.87	53.42	61.94	46.82	48.20	67.80
% Impact (<i>Q5</i>)	7.3	5.2	14.9	14.5	112.9	4.0	17.2	3.6	4.1	10.0
Lower value (< <i>Q5</i>)	0.01 (0.02)	-0.07 (0.05)	1.71 (11.48)	4.78 (5.13)	-0.03 (2.64)	-0.18 (1.15)	0.25 (3.08)	2.10 (1.79)	-0.46 (1.32)	-1.48 (2.32)
Panel B: small firms										
High value (<i>Q5</i>)	-0.10 (0.08)	0.28 (0.15)	19.59 (25.06)	16.56 (7.68)	8.82 (4.03)	5.15 (3.08)	13.28 (5.18)	7.37 (4.41)	2.99 (3.69)	7.52 (6.68)
Mean of outcome (<i>Q5</i>)	0.65	1.64	312.10	127.20	10.47	51.39	63.64	56.43	42.35	64.55
% Impact (<i>Q5</i>)	-15.8	16.8	6.3	13.0	84.2	10.0	20.9	13.1	7.1	11.7
Lower value (< <i>Q5</i>)	-0.02 (0.02)	0.13 (0.06)	-22.21 (14.07)	-3.57 (5.57)	-2.93 (1.95)	1.81 (1.26)	-0.89 (2.36)	0.49 (2.35)	1.48 (0.98)	6.88 (2.23)
Observations	155,646	103,437	103,437	103,437	103,437	103,437	103,437	103,437	107,789	99,558

Notes: This table reports difference-in-differences estimates of the effect of initial patent allowances interacted with large and small firm size on firm and worker outcomes, separately for high and low ex-ante valuable patent applications, in our analysis sample. Large firms are those above median baseline employment, and small firms are those below median baseline employment. Estimates correspond to coefficients on interactions of the designated value category with a post-decision indicator, an indicator for the application being initially allowed, and whether the employment of the firm was above or below median baseline employment. Controls include main effect of value category interacted with a post-decision indicator, firm fixed effects, and art unit by application year by calendar year fixed effects. Standard errors (reported in parentheses) are two-way clustered by (1) art unit, and (2) application year by decision year. Panel A reports the estimates for the large firm indicator interacted with a post-decision indicator and an initial allowance indicator. Panel B reports the estimates for the small firm interacted with a post-decision indicator and an initial allowance indicator. EBITD is earnings before interest, taxes, and deductions. Surplus is EBITD + wage bill. Stayers are defined as those who were employed by the same firm in the year of application. Labor compensation measures total deductions for labor expenses claimed by the firm. “W2 + 1099” measures to the sum of W2 and 1099 earnings divided by the sum of the number of W2’s and 1099’s filed. “% Impact” reports the percent change in the outcome at the mean for winning an initial allowance. Revenue, value added, EBITD, wage bill, surplus, earnings, labor compensation, and W2 + 1099 pay are reported in thousands of 2014 USD.

Table C.4: Within-Firm Inequality

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Avg male earnings	Avg female earnings	Gender earnings gap	Avg inventor earnings	Avg earnings of non-inventors	Inventor earnings gap	Wage bill per worker (Q_1)	Wage bill per worker (Q_4)	Wage bill per worker ($Q_4 - Q_1$)
High value (Q_5)	5.88 (1.97)	0.15 (1.51)	6.90 (1.97)	16.87 (8.47)	2.24 (1.43)	14.92 (7.75)	-0.04 (1.01)	8.06 (2.85)	8.12 (2.56)
Mean of outcome (Q_5)	60.97	36.69	24.42	130.20	47.64	81.49	16.61	110.70	94.17
% Impact (Q_5)	9.7	0.4	28.3	13.0	4.7	18.3	-0.3	7.3	8.6
Lower value ($< Q_5$)	0.35 (1.16)	-0.49 (0.51)	-0.06 (1.06)	-1.24 (4.60)	0.47 (0.81)	-1.72 (4.70)	0.05 (0.32)	2.82 (2.42)	2.77 (2.34)
Observations	95,004	84,562	80,222	52,471	100,901	50,045	82,750	81,536	81,536

Notes: This table reports difference-in-differences estimates of the effect of initial patent allowances on within-firm inequality measures, separately for high and low ex-ante valuable patent applications, in our analysis sample. Estimates correspond to coefficients on interactions of the designated value category with a post-decision indicator and an indicator for the application being initially allowed. Controls include main effect of value category interacted with a post-decision indicator, firm fixed effects, and art unit by application year fixed effects, as in equation (8). Standard errors (reported in parentheses) are two-way clustered by (1) art unit, and (2) application year by decision year. “Gender earnings gap” measures the difference between average male and female earnings at firms where both genders are present. “Inventor earnings gap” measures the difference between average inventor and non-inventor earnings at firms where both inventors and non-inventors are present. “Wage bill per worker (Q_1)” measures the average wage bill in within-firm wage quartile one. “Wage bill per worker (Q_4)” measures the average wage bill in within-firm wage quartile four. “Wage bill per worker ($Q_4 - Q_1$)” measures the difference between average Q_4 and Q_1 earnings. Stayers are defined as those who were employed by the same firm in the year of application. Entrants are defined as those employees who were not employed at the firm in the previous year. Earnings and wage bill are measured in thousands of 2014 USD.

Table C.5: Within-Firm Inequality (Using a Balanced Sample)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Avg male earnings	Avg female earnings	Gender earnings gap	Avg inventor earnings	Avg earnings of non-inventors	Inventor earnings gap	Wage bill per worker (Q_1)	Wage bill per worker (Q_4)	Wage bill per worker ($Q_4 - Q_1$)
High value (Q_5)	7.11 (1.79)	0.21 (1.54)	6.90 (1.97)	14.44 (7.59)	-0.48 (1.86)	14.92 (7.75)	-0.06 (1.01)	8.06 (2.85)	8.12 (2.56)
Mean of outcome (Q_5)	61.45	37.03	24.42	132.30	50.76	81.49	16.55	110.70	94.17
% Impact (Q_5)	11.6	0.6	28.3	10.9	-1.0	18.3	-0.3	7.3	8.6
Lower value ($< Q_5$)	-0.23 (1.12)	-0.18 (0.45)	-0.06 (1.06)	-1.35 (4.91)	0.37 (1.18)	-1.72 (4.70)	0.05 (0.32)	2.82 (2.42)	2.77 (2.34)
Observations	80,222	80,222	80,222	50,045	50,045	50,045	81,536	81,536	81,536

Notes: This table replicates Table C.4 using a common sample for each comparison group. Each column reports difference-in-differences estimates of the effect of initial patent allowances on within-firm inequality measures, separately for high and low ex-ante valuable patent applications, in a subsample of our analysis sample where the composition of firms is held constant across the related outcome columns. Estimates correspond to coefficients on interactions of the designated value category with a post-decision indicator and an indicator for the application being initially allowed. Controls include main effect of value category interacted with a post-decision indicator, firm fixed effects, and art unit by application year by calendar year fixed effects. Standard errors (reported in parentheses) are two-way clustered by (1) art unit, and (2) application year by decision year. “Gender earnings gap” measures the difference between average male and female earnings at firms where both genders are present. “Inventor earnings gap” measures the difference between average inventor and non-inventor earnings at firms where both inventors and non-inventors are present. “Wage bill per worker (Q_1)” measures the average wage bill in within-firm wage quartile one. “Wage bill per worker (Q_4)” measures the average wage bill in within-firm wage quartile four. “Wage bill per worker ($Q_4 - Q_1$)” measures the difference between average Q_4 and Q_1 earnings. “% Impact” reports the percent change in the outcome at the mean for winning an initial allowance. Earnings and wage bill are measured in thousands of 2014 USD.

Table C.6: Earnings Impacts by Stayer Subgroups

	(1)	(2)	(3)	(4)	(5)	(6)
	Avg male stayer earnings	Avg female stayer earnings	Stayer gender earnings gap	Avg inventor stayer earnings	Avg non-inventor stayer earnings	Stayer inventor earnings gap
High value ($Q5$)	13.83 (2.74)	2.73 (1.95)	8.89 (3.51)	17.41 (11.21)	5.75 (1.72)	9.27 (8.35)
Mean of outcome ($Q5$)	85.03	48.35	37.77	156.00	64.37	91.85
% Impact ($Q5$)	16.3	5.7	23.5	11.2	8.9	10.1
Lower value ($< Q5$)	2.17 (1.96)	0.70 (0.83)	-0.83 (1.80)	0.58 (6.31)	2.03 (1.26)	-2.78 (6.91)
Observations	88,100	71,591	66,270	47,063	94,909	42,640

Notes: This table reports difference-in-differences estimates of the effect of initial patent allowances on within-firm inequality measures, separately for high and low ex-ante valuable patent applications, in a subsample of our analysis sample where the composition of firms is held constant across the related outcome columns. Estimates correspond to coefficients on interactions of the designated value category with a post-decision indicator and an indicator for the application being initially allowed. Controls include main effect of value category interacted with a post-decision indicator, firm fixed effects, and art unit by application year by decision year. “Stayer gender earnings gap” measures the difference between average male stayer and female stayer earnings at firms where both genders are present. “Stayer inventor earnings gap” measures the difference between average inventor stayer and non-inventor stayer earnings at firms where both inventors and non-inventors are present. “% Impact” reports the percent change in the outcome at the mean for winning an initial allowance. Stayers are defined as those who were employed by the same firm in the year of application. Earnings are measured in thousands of 2014 USD.

Table C.7: Earnings of Officers / Owners

	All firms				Pass-through entities				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Officer earnings per W2	Wages and salaries per W2	Non-officer comp per W2	Officer earnings per W2	Wages and salaries per W2	Non-officer comp per W2	Avg earnings of owner-employees	Avg pay of owner-employees	Avg W2 earnings of non-owner-employees
High value (<i>Q5</i>)	3.81 (1.31)	0.13 (1.89)	-0.05 (2.10)	7.17 (3.27)	-1.43 (1.66)	-1.67 (1.56)	42.87 (17.01)	84.08 (41.36)	5.99 (2.41)
Mean of outcome (<i>Q5</i>)	15.89	35.07	38.91	14.86	27.85	31.25	149.70	246.60	41.76
% Impact (<i>Q5</i>)	24.0	0.4	-0.1	48.3	-5.1	-5.4	28.6	34.1	14.4
Lower value (< <i>Q5</i>)	-0.83 (0.94)	1.46 (1.01)	1.97 (1.02)	-1.28 (1.38)	0.46 (1.31)	1.22 (1.49)	-0.28 (5.45)	-3.71 (18.72)	0.66 (0.84)
Observations	103,437	103,437	103,437	49,943	49,943	49,943	31,962	45,318	43,531

Notes: This table reports difference-in-differences estimates of the effect of initial patent allowances on officer and non-officer earnings measures for all firms and pass-through entities. Estimates correspond to coefficients on interactions of the designated value category with a post-decision indicator and an indicator for the application initially allowed. Controls include main effect of value category interacted with a post-decision indicator, firm fixed effects, and art unit by application year by calendar year fixed effects. Standard errors (reported in parentheses) are two-way clustered by (1) art unit, and (2) application year by decision year. “% Impact” reports the percent change in the outcome at the mean for winning an initial allowance. Earnings are measured in thousands of 2014 USD. The outcome variables are defined as follows. (1) Officer earnings per W2 is the officer compensation component of labor compensation that we define in A.3. Specifically, the numerator of (1) is Line 12 on C-corporation form 1120, line 7 on S-corporation form 1120S, and line 10 on Partnership form 1065. Each of these line numbers corresponds to the 2010 IRS forms. Similarly, the numerator of Column (2) is the salaries and wages component of labor compensation, which is line 13, 8, and 9 on the C-corporation, S-corporation, and partnership forms, respectively. The numerator of Column (3) is labor compensation less officer compensation, so it includes the other components of labor compensation (i.e., it includes salaries and wages, pension profit-sharing plans, and employee benefit programs). On form 1120, for example, these are lines 13, 23, and 24. Columns (4)-(6) restrict the sample to only S-corporations and Partnerships, but have the same definitions, respectively, as Columns (1)-(3). Column (7) is the average W2 wage earnings that go to firm owners, which was constructed using matched firm-owner data from Smith et al. (2017). Specifically, for each S-corporation and partnership, we calculate the sum of W2 wages that accrue to its owners and divide this sum by the number of owners who also get positive W2 income from the firm. Column (8) is average owner pay, which is the sum of W2 wage earnings that accrue to owners plus their business income divided by the number of owners. Each owners’ business income comes from Smith et al. (2017)’s linkage between firms and owners, which was constructed using Schedule K1 of form 1120S for S-corporations and form 1065 for Partnerships. Being able to identify business owners directly is only possible for S-corporations and partnerships, so only aggregate firm-level business income is available for C-corporations (on line 28 of form 1120). Finally, Column (9) is the total W2 earnings for non-owners divided by the number of non-owner W2s.

Table C.8: Pass-Through Estimates: Three-Year Average of Surplus per Worker

	Wage bill per worker	Avg male earnings	Avg non-inventor earnings	Avg stayer earnings	Avg earnings of stayers minus earnings in app yr	Avg non-inventor stayer earnings
	(1a)	(2a)	(3a)	(4a)	(5a)	(6a)
	OLS	OLS	OLS	OLS	OLS	OLS
Surplus / worker	0.20	0.23	0.16	0.26	0.25	0.21
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Elasticity	0.243	0.242	0.215	0.253	0.224	0.224
	(0.15)	(0.25)	(0.14)	(0.34)	(0.32)	(0.28)
Observations	83,212	77,066	81,632	81,075	81,075	77,939
1 st stage F	8.00	6.21	6.38	8.58	8.58	5.97
Exogeneity	0.284	0.045	0.257	0.179	0.216	0.072

Notes: This table reports OLS and IV estimates of the effect of increases in surplus per worker on selected earnings outcomes using three-year averages of surplus per worker. The excluded instrument is the interaction of top quintile of ex-ante value ξ category with a post-decision indicator and an indicator for the application being initially allowed. Controls include main effect of value category interacted with post-decision indicator and interaction of lower quintile value category with a post-decision indicator times an indicator for initially allowed, firm fixed effects, and art unit by application year by calendar year fixed effects. Standard errors (reported in parentheses) are two-way clustered by (1) art unit, and (2) application year by decision year. “Exogeneity” reports p-value for test of null hypothesis that IV and OLS estimators have same probability limit. Stayers are defined as those who were employed by the same firm in the year of application. Surplus is EBITD (earnings before interest, tax, and depreciation) + wage bill. Earnings, wage bill, and surplus are measured in thousands of 2014 USD.

Table C.9: Sensitivity analysis of retention estimates and model based interpretation

	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline	Low incumbent premium	High retention elasticity	Low retention elasticity	High pass-through rate	Low pass-through rate
Panel A. Calibrated wage premium						
Calibrated inputs						
w_j^l/w_j^m	1.8	1.2	1.8	1.8	1.8	1.8
$d \ln G(w_j^l) / d \ln w_j^l$	1.2	1.2	1.8	0.6	1.2	1.2
π	0.61	0.61	0.61	0.61	0.91	0.31
Model-based outputs						
η	2.7	7.3	4.0	1.4	2.7	2.7
$c'(N_j/I_j)/w_j^m$	1.1	0.2	1.0	1.4	1.1	1.1
θ	0.73	0.88	0.80	0.59	0.73	0.73
ε	6.0	3.3	4.2	-	-	1.7
Panel B. Calibrated elasticity of product demand						
Calibrated inputs						
ε			6.0	6.0	6.0	6.0
$d \ln G(w_j^l) / d \ln w_j^l$			1.8	0.6	1.2	1.2
π			0.61	0.61	0.91	0.31
Model-based outputs						
w_j^l/w_j^m			2.9	1.3	-	-
η			2.7	2.7	-	0.6
$c'(N_j/I_j)/w_j^m$			2.6	0.4	-	-
θ			0.73	0.73	-	0.37

Notes: This table reports estimates of model-based parameters under different calibrations of the retention elasticity and pass-through rate. Panel A allows for different calibrations of the incumbent wage premium w_j^l/w_j^m . Column (1) reports our baseline calibration, where the wage premium is the ratio of the average earnings of incumbent workers to average earnings of recent entrants in the year of application. The baseline retention elasticity $d \ln G(w_j^l) / d \ln w_j^l$ is estimated in Table C.1 and the pass-through rate is estimated in Table 8. Column (2) recalibrates all model-based parameters assuming a lower wage premium of 1.2. Columns (3) and (4) recalibrate assuming the retention elasticity is one standard deviation above and below the baseline, respectively. Columns (5) and (6) re-estimate the model-based parameters assuming the pass-through rate is one standard deviation above and below the baseline, respectively. Panel B holds the baseline elasticity of product demand ε constant and allows the wage premium to vary. Column (1) in Panel B would replicate the baseline results, so it was excluded from this table. Column (2) is not reported in Panel B as it reports the results of calibrating the incumbent wage premium. Dashes denote cases in which the parameter value lies outside a feasible range. See Section 2 for how the parameters are calculated.