

# Retention Effects of Employee Stock Options: Evidence from Bunching at Vesting Dates

Qing Gong, James Liang, Hong Zhang, and Li-An Zhou\*

October 14, 2017

---

\*We thank Nicholas Bloom, François Gerard, Wojciech Kopczuk, Edward Lazear, Paul Oyer, John Shoven, and Petra Todd for many helpful comments and suggestions. We also thank seminar participants at the Beyster Fellowship Symposium for Employee Ownership and Profit Sharing and the Columbia University Applied Micro Research Methods Colloquium. Hong Zhang gratefully acknowledge the financial support from the National Natural Science Foundation of China (Grant No. 71703054). All remaining errors are our own. Gong: Department of Economics, University of Pennsylvania (e-mail: qinggong@sas.upenn.edu). Liang: Guanghua School of Management, Peking University (e-mail: jzliang@gsm.pku.edu.cn). Zhang: Institute of Industrial Economics, Jinan University (e-mail: hongzh@jnu.edu.cn). Zhou: Guanghua School of Management, Peking University (e-mail: zhoul@gs.m.pku.edu.cn). PRELIMINARY AND INCOMPLETE. PLEASE CONTACT THE AUTHORS BEFORE CITING.

## Abstract

Whether employee stock options bring the firm enough benefit to justify the cost is at the center of much debate. We study the retention effect of stock options to reconcile their popularity and high granting costs. We use the bunching design to address the endogeneity problem and identify the causal effects of options on retention. Using a novel administrative panel dataset at the individual employee level, we find option owners delay quitting until options vest to minimize the opportunity cost of quitting. The bunching at the vesting dates is significant and sizable: the quitting of option owners more than doubles shortly after options vest. We find the bunching effect is robust to alternative explanations, stronger when option values are high, unique to voluntary quitters but absent among owners who leave involuntarily. We also verify that the retained option owners have superior performance, and conduct a simple benefit-cost analysis with the baseline estimates. We estimate that the retention benefits contribute greatly to the total benefit of options, which exceeds the granting cost by 95-275%. Accounting for the retention effects of options can avoid the potential underestimation of the benefit-cost ratio, thereby helping to reconcile the popularity of options and the high granting costs.

*JEL classification:* G30; J33; M50; M52

*Keywords:* Stock options; Employee turnover; Bunching

# 1 Introduction

Employee stock options are popular in compensation packages. About 8% of private-sector employees in the United States held stock option ownership plans in 2014, which is already down from the peak of over 13% in the early 2000s.<sup>1</sup> Such popularity is not unique to the U.S. 8-16% of Chinese firms offered employee stock options in 2012, despite a small and highly regulated derivatives market.<sup>2</sup> But broad-based employee stock options are costly to grant. And there is still much debate about whether options bring firms enough benefits to justify the cost, and, if not, how to reconcile their popularity with the benefit-cost discrepancy.

One plausible justification of granting stock options is they align the interests of employees with those of the firm. Hence option recipients will exert more effort and increase firm value. But the incentivizing effects can be limited by free-riding in multiple-agent settings like most modern firms (Lazear, 2004). Previous empirical studies show this is indeed the case (Barron and Waddell, 2003; Oyer and Schaefer, 2005). Moreover, the incentive hypothesis is hardly sufficient to explain the popularity of options among lower-level employees. These “rank-and-file” employees have little impact on firm value, are often risk-averse, and hold ill-diversified portfolios. So they tend to undervalue the options, which makes granting options less cost-effective than other forms of compensation (Hall and Murphy, 2003).

The limited benefits from incentivizing employees and the popularity of options can be reconciled in two ways. One is to discard the optimality assumption of firm behavior, and see options as a mistake. This view is articulated by Hall and Murphy (2003), who conjecture that options are popular because firms underestimate the costs. Alternatively, one could explore other benefits of options. For example, Oyer and Schaefer (2006) and Lazear (2004) propose that options help the firm attract and retain employees at all levels. The options can also sort employees and select those who are optimistic about the firm’s prospects. Micro-level evidence of these alternative explanations, however, is currently

---

<sup>1</sup><http://www.nceo.org/articles/widespread-employee-ownership-us>

<sup>2</sup><http://www.china-briefing.com/news/2011/05/12/using-stock-options-to-attract-and-retain-key-employees-in-china.html>

tenuous at best.

In this paper, we study the retention benefits of employee stock options. We test and quantify the extent to which options reduce turnover, and weigh the retention benefits against the cost of granting. A more comprehensive assessment of the benefits of employee stock options is important for firms and policy makers alike, the latter of whom often provide tax benefits to encourage employee option ownership (Dube and Freeman, 2010).

To measure the retention effects, we need to draw causal conclusions about how employee turnover react to options. One challenge for identification is the endogeneity of option ownership. Option owners are often a highly selected subgroup of employees, who might differ from their non-owner peers in quitting decisions even without options. A direct comparison of the turnover rates between owners and non-owners is thus confounded by unobserved employee heterogeneity.

We resolve the endogeneity problem using the bunching design on a unique dataset. First, we use detailed longitudinal data with rich information on individual employee characteristics, wages, option ownership and transactions, and other personnel records from a leading Chinese online travel agency. Second, we take advantage of the predetermined vesting schedule of the firm's employee stock options. The options are granted throughout the calendar years, and vest in three equal batches every 12 months after an initial lock-in period. When an option owner quits, she has to give up all the unvested shares, but can still exercise vested shares within 90 days of quitting. Because the number of options forgone drops by one-third on the vesting day, the option owner's quitting cost also decreases discontinuously. This generates bunching of option owner turnover on the vesting dates. We use the bunching design to pin down the retention effects of employee stock options. The causal relationship that we are able to identify adds to the micro-level empirical evidence that is much needed in this literature.

We find significant bunching at the vesting dates. The quitting rate of option owners increases by 0.5 percentage points immediately after options vest. The excess quitting is large relative to the average monthly quitting rate of 0.3% before vesting dates. We use option transaction data to show that the bunching is indeed driven by owners' strategic delay in order to minimize the cost of foregone unvested shares. We also find that the bunching

is robust to tests of alternative explanations, such as cyclicalities within the calendar year or the contract year. Moreover, we use option owners who leave involuntarily to show that bunching is unique to voluntary quitters, i.e. those who can choose when to leave the firm. These robustness and falsification tests corroborate our baseline finding. They also indicate that the bunching at vesting dates captures the retention effect of options instead of confounding factors.

We extend the empirical analysis along two dimensions. In the first extension, we look into option owners who quit shortly *before* vesting dates. We find that quitting “too early” seems to be the result of better outside offers rather than optimization frictions. In the other extension, we do a quick benefit-cost analysis. We first show that employees who are retained by options have superior performance, which is the premise for discussing any retention *benefits*. We then weigh the cost of granting options against two possible benefits of options: the wage savings and the retention benefits. We find that the retention benefits contribute greatly to the total benefit of option granting, which more than justifies the granting costs. The overall benefit-cost ratios range from 1.95-3.75. Not accounting for the retention effects would lead to considerable underestimation of the benefit-cost ratio. Hence accounting for the retention effects of options can avoid the potential underestimation of the benefit-cost ratio, thereby helping to reconcile the popularity of options and the high granting costs.

Our study is related to a large body of literature on profit sharing in general, and employee stock options in particular. Kruse (1993) examines a panel of 500 U.S. firms over 21 years and argue that profit sharing plans help boost productivity and improve economic stability by lowering unemployment rate. Lazear (2000) also shows that variable payment schedules can partially resolve asymmetric information about efforts, and work as an incentive provision tool. The popularity of options leads researchers to study why employers grant stock options and whether the options work as expected. Earlier studies usually adopt the incentivization explanation and focus on firm executives. Aggarwal and Samwick (1999) find that the more sensitive an option-holding executive’s pay is to performance, the better the firm performs. But the evidence still cannot explain the granting of options in the first place. Core and Guay (2001) find that firms with lower monitoring

costs have smaller fractions of option-holding employees, which sheds light on why firms award options. Other studies usually find no discernible relationship between employee options and firm value (e.g. Kedia and Mozumdar (2002)).

Because the incentivization hypothesis cannot fully justify the popularity of options, researchers begin to explore other channels. Lazear (2004) provides theoretical foundations for the selection, or sorting, effects of options. Oyer (2004) proposes a model in which a worker's outside option correlates with firm value and stock prices. Hence firms grant options as an alternative to wage adjustment and turnover, both of which are costly. Then Oyer and Schaefer (2005) evaluate the three alternative explanations for granting options—*incentivization, sorting, and retention*. Their empirical work draw upon a combination of firm-level option grant data and survey data on salary and option packages of middle-level executives, which reject the incentivization explanation and find the latter two more plausible. Oyer and Schaefer (2006) use the above-noted survey data to to conduct a detailed estimation of firm's cost of granting options. They find that accounting concerns are unlikely to be the driving force behind option granting decisions, and that the empirical patterns are more consistent with the retention justification. Carter and Lynch (2004) study option repricing and find it to be negatively correlated with overall employee turnover but not executive turnover. Dube and Freeman (2010) use survey data and find that equity compensation and shared decision-making arrangements are jointly associated with lower turnover.

Our paper contributes to the existing literature in two ways. First, we provide new empirical evidence on the effect of stock options at the individual employee level. The vast majority of empirical studies look at firm-level responses, which is largely due to data limitations. Existing papers on employee-level effects almost all use data on company executives (e.g. Aggarwal and Samwick (1999); Mehran and Tracy (2001); Barron and Wad-dell (2003)). Using a novel administrative dataset on employees at all levels within a large firm, we are able to closely examine the effect of broad-based stock options on employee turnover. To the best of our knowledge, the only other paper using detailed employee-side data in the literature is Cowgill and Zitzewitz (2014), which studies the incentivizing effect

of restricted stocks on Google employees.<sup>3</sup>

Our second contribution is to identify the causal effect of option ownership on retaining employees with the bunching approach. The endogeneity problem often plagues empirical studies in the literature: option owners and non-owners are likely to have unobserved differences that affect turnover and performance. Hence direct comparisons between the two groups are usually subject to bias. The bunching design addresses the endogeneity problem and identifies the retention effects of options. It is also known for the robustness and strong internal validity.<sup>4</sup>

The rest of the paper is organized as follows. We introduce in Section 2 the empirical background and the data. We discuss the vesting schedules and evolving values of options, and how they could affect turnover. We also highlight option owner heterogeneities, which guide the design of our empirical analysis that follows. We describe in Section 3 the bunching framework and the joint estimation of the counterfactual distribution and the affected range. We discuss at the end of this section the advantages of bunching and how it differs from the closely related regression discontinuity design. Then we present the baseline estimates of bunching in Section 4, followed by various robustness and falsification tests that corroborate the key findings. We extend the empirical analysis in Section 5 to examine the suboptimal turnover responses by some option owners. We also explore in this section the benefits and costs of employee options. We conclude in Section 6.

---

<sup>3</sup>Another loosely related paper using data on lower-level employees is Huddart and Lang (1996). The study takes an accounting perspective and examines the implications of option exercising on corporate debt and regulation.

<sup>4</sup>Cowgill and Zitzewitz (2014) is another paper that examines the causal effects. Taking advantage of Google's unique equity pricing policy, they identify the incentivizing effects of equity compensation from exogenous variations in stock exposure. They find the incentives to be fairly weak. We take a different approach, using the discontinuous change in the opportunity cost of quitting to deal with unobserved employee heterogeneity. We also focus on the retention effects on employees who plan to quit, instead of the effects on employees who choose to stay.

## 2 Empirical Background and Data

### 2.1 Firm background and option grants

Our longitudinal administrative data comes from a leading online travel agency in China. The firm was founded in the late 1990s, listed in NASDAQ in early 2000s, and has over 20,000 employees as of now. The firm categorizes the job positions vertically into 10 levels.<sup>5</sup> Newly hired employees can start from any level depending on their profile upon entry, and then move to higher levels over time. Levels 1 to 4 are junior employees, senior employees, heads of teams, and senior heads of teams, respectively. They receive options as an award for excellence in work performance. Around 2% of levels 1-4 employees (1,075 individuals) have received option awards. Levels 5-10 are managers, senior managers, and higher-level positions all the way to vice president. Almost all of them (93.2%, or 718 individuals) receive stock options as a part of the compensation package.

The firm grants employee stock options at least twice a year since 2000. There are two types of plans: Plan 1 options are offered from 2000 to 2009, and Plan 2 from 2008 onward. The two types of option plans differ, among other things, in the vesting schedule. Plan 1 options vest in three equal batches every 12 months after an initial 12-month lock-in period. Suppose an employee receives options on July 20th, 2005, which we refer to as the beginning of month 0. Then one-third of her total granted options vest at the beginning of month 12 on July 20, 2006. An additional one-third vest at the beginning of month 24 on July 20, 2007. And the final one-third vest at the beginning of month 36 on July 20, 2008. Plan 2 options also vest in three equal batches, but the initial lock-in period is 24 months instead of 12. So the three vesting dates are at the beginnings of months 24, 36, and 48.

An option owner has to give up her *unvested* shares upon leaving the firm, but can keep and exercise the *vested* shares for 90 days. Any vested shares unexercised thereafter revert to the company, too. As a result, the value of foregone options affects the cost of quitting. And option owners who consider quitting will take the vesting dates into account, especially when the value of options is high.

---

<sup>5</sup>Top executives (e.g. the chief executive officer and chief financial officer) are technically at Level 11. But we exclude them from our sample because their option plans and turnovers differ substantially from other employees, who are our focus in studying broad-based employee stock options.



Another notable difference between the two types of option plans is the strike prices. Plan 1 options have significantly lower strike prices than Plan 2 options. The difference in option values motivates our separate examination of Plan 1 and Plan 2 options later in the paper.

[FIGURE 1 ABOUT HERE]

Figure 1 illustrates the vesting schedules. The horizontal axes in all three panels are months after the grant date. The top panel plots the fraction of vested Plan 1 (solid) and Plan 2 (dashed) options over time. Plan 2 options have a longer lock-in period and all shares remain unvested for the first 24 months. The middle panel plots the average *expected* value of vested options, which is the Black-Scholes value of the vested fraction calculated on the granting date. The bottom panel plots the *actual* Black-Scholes value of vested options.

All three panels show that vested options increase discontinuously on vesting dates, both in terms of shares and values. Put differently, the opportunity cost of quitting drops discontinuously on vesting dates.<sup>6</sup> We will use the discontinuous change in quitting costs to design the bunching approach.

## 2.2 Data and sample construction

The dataset we use spans from 1999 to 2014, covering all employees ever joined the firm by the end of 2012 (over 50,000 unique individuals). It contains a rich set of information that the firm's human resources department keeps track of. In particular, it documents detailed option ownership and transactions, monthly wages, promotion records, and other employee characteristics.<sup>7</sup>

We restrict the sample to option-owners in the baseline analyses. Because some employees receive option grants more than once, we restructure the employee-month level

---

<sup>6</sup>Although the bottom panel shows an increase in Plan 1 option values in month 48 (after all options vest), it does not invalidate the uniqueness of discontinuity on vesting dates. Also note that the change around month 48 is continuous. In fact, it is no more drastic than the change in option values in other non-vesting months.

<sup>7</sup>We convert all money values to real 2014 RMB yuan using province-level urban CPIs.

data so that each observation is a unique option grant, employee, and month combination. The restructured data can distinguish between batches of options for the same individual, which is important because the batches have different vesting dates and values.

For each batch of options granted, we observe the grant date, the vesting dates, the strike price, the evolving spot market prices, and the number of vested and unvested shares at each point in time. For each option owner, we observe her job level, department, promotions, itemized monthly wages, and the date and cause of leaving the firm if separation happens.<sup>8</sup> We also observe the option owner's characteristics such as age, gender, education, and work experience. In the end, we get an unbalanced panel of 185,593 observations from 2,482 distinct option-employee combinations.

[TABLE 1 ABOUT HERE]

Table 1 reports the summary statistics on option-employee characteristics. On average, option-owners of Plan 1 receive 9,123 shares, and those of Plan 2 receive 7,079 shares. The number of shares varies substantially across individuals, though. In relative terms, the value of Plan 1 options is about 14.53 times the annual base wage of their owners, and that of Plan 2 is 5.69 times. The large difference in the two plans' relative values is partially explained by the difference in owner wages: the average monthly wage of Plan 1 owners is slightly lower than Plan 2 owners, which is consistent with Plan 1 options being granted in earlier years.

As for the equity portfolio held by option owners, more than 60% of option owners have other options on hand when they receive the current batch of options. The majority own other Plan 1 options, which is consistent with the longer history of Plan 1.

The distribution of owner job levels follows an inverse U-shaped pattern for both plans. Fewer owners are at very low levels when receiving the options, more in intermediate levels, and fewer in very high levels. The initial increase is driven by better chances of receiving options when an employee is at higher levels. And the decrease beyond level

---

<sup>8</sup>In the empirical analyses that follow, we define monthly "wage" as the monthly average of an employee's total annual cash compensation, which include wages/salaries, cash bonuses, and all other cash benefits. We use this comprehensive measure instead of pure wage because the total compensation may be more relevant to the employees' decision-making that we are interested in.

5 is because of the drastic decline in the number of high level positions. The distribution for Plan 1 owners is slightly skewed toward the lower end, with a larger fraction of level-3 and level-4 owners.

The statistics reveal the widespread heterogeneities across option plans, both for the options and for the owners. The heterogeneities call for separate examinations of the two plans when measuring the retention effects. Moreover, option owner heterogeneity re-emphasizes the need for a robust identification strategy to pin down any causal relationship.

### 3 Empirical Strategy

#### 3.1 Option vesting and the bunching design

Now we formalize the bunching design that identifies the retention effect of employee options. We first define time  $t$  for each option-employee combination as months elapsed since the grant date. Take the Plan 1 owner who receives her options on July 20, 2005 as an example. For this option-employee combination,  $t = 0$  from July 20 to August 19, 2005;  $t = 1$  from August 20 to September 19, 2005; and so forth. The first one-third of her options vest on the first day of  $t = 12$  on July 20, 2006. The next one-third vest on the first day of  $t = 24$  on July 20, 2007. The last one-third vest on the first day of  $t = 36$  on July 20, 2008. We introduce the relative notion of time,  $t$ , because it is the running variable in the bunching framework. And the predetermined option vesting dates are the treatment-determining thresholds. Treatment in our empirical setting is the vesting of options, which reduces the cost of quitting discontinuously.

We use the bunching design to study how options change option owners' quitting decisions. The value of unvested options foregone due to quitting decreases discontinuously on the vesting date. But once after the vesting date, the marginal cost to delay quitting remains unchanged, although the average cost does decrease. So option owners are more likely to quit immediately after vesting dates, and less likely to quit before vesting dates. The excess mass in the distribution of quitting after vesting dates gives rise to bunching with a "notch." Correspondingly, the missing mass before vesting dates forms a "hole." The excess mass in the notch and missing mass in the hole will exactly offset each other assuming no uncertainty and no optimization frictions.

The counterfactual distribution and the affected range are the two pillars of estimating bunching effects. Here the counterfactual distribution is that of option owner quitting in the absence of options. It is the benchmark for quantifying missing or excess mass around the vesting dates. The standard approach is to first exclude the data points in months affected by the bunching (i.e. those in the *affected range*); then fit a flexible function, often a polynomial, to the remaining sample; and finally interpolate over the previously excluded affected range to get the entire distribution (Kleven, 2016).

The affected range for each vesting date is the period of time over which option vesting affects owner turnover. Note that each option-employee combination is subject to three potential vesting dates. Therefore we define treatment *locally* as being affected of a nearby vesting date that is at most six months away: the above-noted Plan 1 option owner in September 2006 might be affected by the treatment effects of the first vesting date (July 20, 2006). But we rule out the possibility that her quitting decision is affected by the second vesting date (July 20, 2007) that is about 10 months ahead. Conceptually, the assumption ensures that all option-employee combinations are subject to at most one treatment at any point in time. Empirically, it avoids overlap between the affected ranges of two adjacent vesting dates.

We do not impose further assumptions on the affected ranges. We do not assume the affected ranges are the same across the three vesting dates, or that they are symmetric before and after each vesting date. Instead, we estimate the affected range jointly with the counterfactual distribution, and let the equality between the excess and missing masses govern the affected range.<sup>9</sup>

### 3.2 Joint estimation of counterfactual distributions and affected ranges

We follow Kleven and Waseem (2013) and take an iterative approach to estimate the counterfactual distribution and the affected range jointly. We start with an initial guess of the affected range,  $[z_-^0, z_+^0]$ , which we take as given for now and run the regression:

$$Q_t = \sum_{j=0}^p \beta_j^0 \cdot (t)^j + \sum_{\tilde{t} \in [z_-^0, z_+^0]} \gamma_{\tilde{t}}^0 \mathbf{1}\{t = \tilde{t}\} + \varepsilon_t \quad (1)$$

where  $Q_t$  is the percentage of employees who quit in period  $t$  and  $p$  is the order of the polynomial.<sup>10</sup> We use  $p = 5$  for all relevant specifications throughout this paper.  $\gamma_{\tilde{t}}^0$  captures

<sup>9</sup>The equality between missing mass before the threshold and the excess mass after it requires the absence of extensive margin responses, which we discuss in Section 3.2. Apart from the missing and the excess masses, the distribution farther away might also shift because of the treatment. But Kleven (2016) shows that the econometrician can safely ignore the potential shift and still see the equality as holding for estimation purposes.

<sup>10</sup>We normalize the vesting dates,  $z^*$ , to  $t = 0$  to improve the performance of linear regressions with high-order polynomials present.

the effect of the notch on option owners' quitting decisions in the affected range  $\tilde{t} \in [z_-^0, z_+^0]$ . The counterfactual distribution of quitting is obtained from the predicted values

$$\hat{Q}_t^0 = \sum_{j=0}^p \hat{\beta}_j^0 \cdot (t)^j \quad (2)$$

where the effects from the affected range,  $\hat{\gamma}_{\tilde{t}}^0$ , are omitted. Given the counterfactual distribution, we calculate the missing mass before the vesting date  $z^*$ :

$$\hat{M}^0 = \sum_{t \in [z_-^0, z^*]} (Q_t - \hat{Q}_t^0) \quad (3)$$

Similarly, we calculate the excess mass (i.e. bunching) after the vesting date:

$$\hat{B}^0 = \sum_{t \in [z^*, z_+^0]} (\hat{Q}_t^0 - Q_t) \quad (4)$$

Assuming no extensive margin response, we should expect the missing to be equal to the excess mass,  $M = B$ . The absence of no extensive margin response implies that option owners only adjust *when* to quit but not *whether* to quit due to an approaching vesting date. Options might also retain employees by making them give up quitting entirely, but the extensive margin effect is not identifiable. This limitation is not unique to the bunching design, though, because of these employees are observationally equivalent to those who never want to quit. For now, we use the equality of  $M$  and  $B$  to adjust the affected range. If  $\hat{M}^0 \neq \hat{B}^0$ , we update the guess for the affected range to  $[z_-^1, z_+^1]$ . We iterate until the distance between  $\hat{M}$  and  $\hat{B}$  is minimized and denote the optimal estimates of distribution parameters as  $\hat{\theta}^* \equiv \left( \{\hat{\beta}_j^*\}_{j=0}^p, \{\hat{\gamma}_{\tilde{t}}^*\}_{\tilde{t} \in [z_-^*, z_+^*]} \right)$  and the bunching effect as  $\hat{B}^*$ .

We then use parametric bootstrapping to calculate the standard errors of bunching,  $\hat{B}^*$ . For each bootstrap sample  $k = 1, 2, \dots, K$ , we follow these steps:

1. Draw with replacement a vector of random "errors"  $\nu^{(k)}$  from the empirical distribution of the residuals from (1);
2. Generate a new vector of  $\hat{Q}^{(k)} = X' \hat{\theta}^* + \nu^{(k)}$ ,  $X$  being the matrix of regressors in (1);

3. Re-estimate (1) on  $\hat{Q}^{(k)}$  to get  $\hat{\theta}^{(k)}$  and  $\hat{B}^{(k)}$ .

Finally, we calculate the sample standard deviations of  $\hat{B}^{(s)}$ , and use it as an estimate of the standard error of  $\hat{B}^*$ :

$$\hat{\sigma}_{\hat{B}^*} = \sqrt{\frac{1}{K-1} \sum_{k=1}^K (\hat{B}^{(k)} - \bar{\hat{B}})^2} \quad (5)$$

where  $\bar{\hat{B}} = \frac{1}{K} \sum_{k=1}^K \hat{B}^{(k)}$  is the average of  $\hat{B}^{(k)}$  across the  $K$  bootstrap samples.

### 3.3 Discussion: advantages of bunching and distinctions from regression discontinuity

We conclude this section with a discussion of the bunching design. We first highlight the advantages of bunching as an identification strategy, and then how the empirical setting suits a bunching framework instead of regression discontinuity (RD) design.

Taking the bunching approach is a natural result of the empirical setting we are examining. It is also advantageous in at least two ways. First, it identifies the causal relationship between stock option ownership and employee retention in the context of a vesting schedule. As Kleven (2016) points out, simply plotting the observed distribution often reveals bunching and causality. We will show the plots at the very beginning of Section 4. Second, the bunching approach explicitly allows for imperfections of agents' decision-making process, such as errors and frictions of optimization. We will exploit this feature and look into the option owners who quit right before option vesting in Section 5.

A seemingly close alternative empirical strategy is the regression discontinuity (RD) design. RD is another powerful tool to identify causal effects with strong internal validity.<sup>11</sup> It is not uncommon to use time as the running variable in RD applications (RD in Time, or RDiT), although Hausman and Rapson (2017) point out several potential fallacies of RDiT. RDiT does not suit our study, however, because of an important distinction from the bunching framework: in RD, the running variable cannot be a choice of decision makers or subject to any other manipulation; but in a bunching framework, the running

<sup>11</sup>See Imbens and Lemieux (2008) for a survey of the literature.

variable is an explicit choice. The running variable in our setting is the time of quitting, hence a choice of option owners. It is exactly the change in option owners' chosen quitting time that pins down the retention effects. Hence we adopt the bunching framework instead of RDiT.

## 4 Estimation Results on Bunching and Retention Effects

### 4.1 Preliminary evidence of bunching

We first plot the observed monthly quitting rates of option owners as preliminary evidence of bunching. We restrict the attention to *quitters*, employees who voluntarily leave the firm. Employees who leave involuntarily (e.g. firing or layoff) are excluded from the main sample.<sup>12</sup>

In Figure 2, the horizontal axes show months after option granting, with the grant date being the first day of month 0. The top and bottom panels plot the fraction of quitters among initial option owners of Plan 1 and Plan 2, respectively. Vesting dates are marked by dotted vertical lines with labels showing the cumulative fraction of options vested so far. The estimated affected ranges are shaded in light gray.

The figure highlights the dip in quitting shortly before the vesting dates and the sharp increase afterwards. For example, the average fraction of quitters per month among Plan 1 owners is about 0.3% in the first 6 months. It drops to merely 0.1% around month 10-11; then jumps to almost 1% in month 12 after the first one-third of options vest. The drop and jump in quitting rates correspond to the missing and excess masses in the bunching design. Moreover, bunching is unique to vesting dates. No discernible bunching exists around months 48 or 60, by which time all Plan 1 options have vested. The bottom panel for Plan 2 owners shows similar patterns, with jumps in quitting rates only in months 24, 36, and 48, but not in earlier or later months.

[FIGURE 2 ABOUT HERE]

---

<sup>12</sup>We will examine involuntary turnover later in this section as a falsification test.



The plots uncover preliminary evidence for the retention effects of options. To reiterate the point by Kleven (2016), plotting the raw data in a bunching framework is already informative of the casual relationship. The bunching effects shown in the figure do differ across plans and vesting dates in magnitude and statistical significance. We defer the discussion of it to the next subsection after the formal estimation.

## 4.2 Baseline estimates of bunching

We present the baseline bunching estimation results in Table 2. Column (1) shows the vesting date whose treatment effect we estimate. Column (2) quantifies the bunching effect, measured by the excess mass in the quitting distribution after the vesting date. Column (3) shows the expected length of retention due to option vesting. Columns (4)-(5) report the estimated affected range.  $z_-^*$  and  $z_+^*$  is the number of affected months before and after the vesting date, respectively. Columns (6)-(8) report the sample size used in the estimation for each threshold. Bootstrapped standard errors are in parentheses. We report the results separately for Plan 1 and Plan 2 in panels A and B, respectively.

[TABLE 2 ABOUT HERE]

For Plan 1 owners, excess quitting after the vesting dates ranges from 0.44 to 0.56 percentage point. The bunching effects are statistically significant for all three vesting dates. Moreover, recall that the average monthly quitting rate of Plan 1 owners is around 0.3% before vesting dates. Hence the bunching effects are also economically significant. The affected ranges remain fairly stable at 1-3 months before and after vesting. Estimates on the affected range before vesting ( $z_-^*$ ) are noisier than those after vesting ( $z_+^*$ ).

For Plan 2 owners, bunching effect sizes and significance are more volatile across vesting dates. The excess quitting after the first vesting date in month 24 is 0.28 percentage point and significant; that after the second vesting date in month 36 is as high as 0.83 percentage point; but that after the last vesting date in month 48 is 0.26 and statistically indistinguishable from 0. One explanation is the lower value of Plan 2 options. Option owners delay quitting to reap the gain of options that will soon vest. So options with

lower values provide less incentive to delay quitting; and options out of the money may provide no incentive at all. That Plan 2 options are granted from 2009 onward exposes them to greater shocks in the 2008-09 financial crisis. We explore this possibility further in the next section.

[FIGURES 3 and 4 ABOUT HERE]

Figures 3 and 4 illustrate the estimates in Table 2. The three plots in each figure correspond to the three vesting dates. Take Panel a in Figure 3 for example. The solid line traces out the observed distribution of quitters. The dashed line is the counterfactual distribution, estimated by fitting a flexible quintic polynomial over months 0-18, excluding the affected range. The resulting counterfactual distribution closely resembles the observed one fairly well, especially given the relatively small number of data points. The vertical dotted lines mark the boundaries of the affected range, which is jointly estimated with the counterfactual distribution. The estimated range spans from month 10 to month 14. The figure highlights the missing mass of quitters in months 10-11; and the excess mass of quitters after option vesting, in particular in month 12.

We have two caveats about the baseline estimates. First, the estimation of counterfactual distributions uses a small number of data points. To fit a *distribution* of quitters, we need to aggregate the large number of option-employee-month level observations to the month level. We also estimate bunching around the three notches separately, which further shrinks the sample size for this step of estimation.

We take two approaches to ameliorate the small sample problem. The first is to pool all months and estimate bunching at the three notches together for each plan. It moderately increases the sample size, though at the cost of imposing stronger assumptions. Most notably, it assumes there is one counterfactual distribution for all months after the grant date. The second approach we take is using weekly turnover rates to estimate the counterfactual distribution, separately for each notch. The sample size is roughly four times as large as that in the baseline specification. But weekly turnover rates are also more volatile, rendering the estimates noisier. We present the results of these two approaches in the Ap-

pendix (Figures A1-A3), and find them to be qualitatively similar to the baseline. Despite some limitations, these alternative estimates provide supplemental evidence of bunching with larger samples, and strengthen the baseline findings.

The other caveat about the estimates is the differential bunching effects around the three vesting dates. The excess mass vary across the notches, especially for Plan 2 options. One reason is that quitting decisions around later vesting dates are complicated by other factors. A few examples are the lack of salience—option owners pay less attention to the second and third vesting dates; non-linear cost of delaying quitting—once option owners postpone quitting beyond the first vesting date, their marginal costs of further delay may be increasing or decreasing. These factors contaminate the bunching estimates at the second and the third vesting dates, but have limited impacts on the first. Hence the bunching found at the first vesting date is the most precise among the three.

Bunching at the first vesting date is also a lower-bound measure of the retention effects of options because of potential extensive margin responses. Some employees who considered quitting before receiving options may decide not to quit at all after becoming option owners. They are observationally equivalent to those who never considered quitting in the absence of options. Hence the extensive margin effects are not captured by the bunching estimates.

### **4.3 Mechanism and robustness analyses**

The baseline estimates show strong and substantial bunching at the vesting dates. Now we use a series of tests to show that the bunching indeed reflects the retention effects of options but not confounding factors that coincide with option vesting. First, we present evidence on the mechanism through which vesting dates affect turnover by looking into the exercising of options. Then we examine how the retention effect varies with the value of options. Next, we show that bunching is not driven by reference points such as certain months within the calendar year or the contract year. We conclude the analyses with a falsification test on the sample of option owners who leave the firm involuntarily and thus have no control over the timing of their exits.

### 4.3.1 Mechanism of bunching: quit timing, option vesting, and exercising

In the bunching framework, option owners delay quitting until their options vest because of a discontinuous drop in the opportunity cost of quitting. This implicitly assumes that option owners would want to exercise the options when they quit. If they have no intention of exercising options, then the vesting date will not affect their decision of when to quit. Now we verify this assumption using data on option transactions.

Figure 5 tracks the percentage of option owners who exercise all vested shares on hand. For each month after the grant date, we calculate the number of option owners who exercise at least the number of shares that have most recently vested, divided by the total number of option owners still in the firm at the time. In the top panel, we find a discontinuous increase in the fraction of Plan 1 owners who exercise all the newly vested shares immediately after the vesting date. It shows that owners do want to exercise the options—in fact, they do so as soon as they can.<sup>13</sup>

[FIGURE 5 ABOUT HERE]

Figure 6 further plots the exercising patterns of option owners who quit. The horizontal axes are months relative to an option owner's exit. Month 0 is the month of quitting; months with negative indices are those before the quit; and months with positive indices are those after the quit. For option owners who quit, we track the cumulative number of options exercised as a fraction of total shares initially granted. We then take the average across quitting owners and plot it against time.

[FIGURE 6 ABOUT HERE]

First note that the exercising patterns are consistent with the firm's policy. The fractions do not reach 100% because some option owners quit while still having unvested options, which revert to the firm immediately. The fractions remain unchanged after month 3 because quitting owners can exercise vested shares only within 90 days after they quit.

---

<sup>13</sup>There is similar evidence for Plan 2 owners shown in the bottom panel, although the pattern is much noisier.

Second, the exercised fraction remains fairly flat for both Plan 1 and Plan 2 owners, until about four months prior to quitting. Then it starts to increase more rapidly, especially one month prior to the exit.

Figures 5 and 6 suggest that option owners tend to exercise vested shares on hand before quitting, and that vesting dates have substantial impacts on option exercising. Hence the driving force of bunching at vesting dates is more likely to be option vesting than other factors.

### 4.3.2 Bunching size and option values

We observe bunching because option owners alter the quitting time to minimize the cost of foregone unvested options. An immediate corollary is that the magnitude of bunching effects would vary with the value of options: to delay quitting until options vest is more profitable when the options are more valuable.

Now we examine the quitting of option owners when their option values are high and low before the vesting date, respectively. High-value options are defined as those whose Black-Scholes value is above the median. Before showing how option values affect bunching, we acknowledge the endogeneity of option values. Employees who receive more valuable options may have unobservable characteristics that affect quitting. But what we are about to compare is the difference in bunching sizes between high- and low-value groups. Since bunching captures the *excess* quitting after vesting dates, it is unlikely to be affected by the *level* of quitting among the two groups (which *is* affected by endogeneity). Hence the endogeneity of option value does not invalidate the comparison.

We run the baseline estimation on the each pair of high- and low-value subsamples separately and report the results in Columns (2)-(3) of Table 3. The bunching effect on Plan 1 owners who hold high-value options is very similar to that on low-value option owners around the first vesting date, but substantially larger around the other vesting dates. The contrast between high- and low-value options is even stronger for Plan 2. Owners with high-value options display greater bunching around all three vesting dates than their colleagues who hold low-value options.

[TABLE 3 ABOUT HERE]

For robustness, we also use an alternative definition of high- and low-value options. We measure the Black-Scholes value of options at the *granting dates*, and report the comparison of bunching effects in Columns (4)-(5) of Table 3. The results are qualitatively similar for Plan 1 owners. For Plan 2 owners, however, the comparison is less clear when we break down the sample by option values upon granting, especially for the last two thresholds. This could be the result of the longer lock-in period of Plan 2 options, which makes the granting date option values less relevant to owners' quitting decisions 3-4 years later.

The heterogeneous effects show that the size of bunching does vary depending on option values. The implication is twofold. First, it corroborates the retention effects of options by establishing the connection between option value and the size of bunching. Second, it explains why we observe smaller and more volatile bunching by Plan 2 option owners, who hold lower-value options that are also subject to more price fluctuation.

### 4.3.3 Potential reference point effects

Next, we test for one particular confounding factor, reference points for quitting. Reference points are times at which employees are systematically more likely to quit. They could arise for various reasons, such as cyclical, institutional rules, or mere social norm. The reference points also generate excess quitting even if there are no financial incentives involved. They could result in the overestimation of treatment effects if they coincide with the vesting dates (Kleven, 2016).

Of particular relevance to our empirical setting is the possibility of quitting at the end of a contract year. Employees usually sign 1-year to 5-year contracts when they join the firm. Suppose they tend to quit after completing full contract years, then contract year ends become a reference point. If the contract year ends also coincide with option grant dates, then the bunching we observe would be a contaminated overestimate for retention effects.

To test whether reference points are driving the results, we first compare the distri-

bution of quitting over the contract year between option owners and non-owners (Figure 7, top panel). Non-owners do tend to quit shortly after completing a contract year. But there is no such tendency for option owners: they are almost equally likely to quit in any month of the contract year. The contrast between option owners and non-owners rejects the hypothesis that bunching is driven by the ending of contract years. In fact, it suggests that the baseline bunching estimates are *underestimated*, because the reference point effect partially offsets the retention effects.

[FIGURE 7 ABOUT HERE]

We also compare the quitting distribution over the calendar year in the bottom panel of Figure 7. Intra-year cyclicalities are essentially nonexistent in the quitting time of non-owners. But some months (notably February, May, June, and September) do see slightly more quitting of option owners.

It is not clear why these months seem to become reference points for employees. So we formally test reference point effects by conditioning the empirical distribution of quitters on fixed effects of months that are potential reference points. We include month fixed effects for February, May, June, and September. To counteract the potential underestimation bias discussed above, we also add fixed effects of the first, second and third month after the end of contract years. Then we estimate the bunching model based on the adjusted empirical distribution.

We report the estimates in Table 4 and illustrate the bunching effects in Figure 8. We find that Plan 1 affected range estimates do not change qualitatively after controlling for potential reference points. Estimates on the excess mass do change slightly in size but remain significant.

[TABLE 4 ABOUT HERE]

[FIGURE 8 ABOUT HERE]

Estimates on Plan 2 owners are similar, except that the bunching effect at the first vesting date (month 24) diminishes. One possibility is that Plan 2 options have lower values

on average, and experience greater fluctuations. When the option values are too low, the latent bunching effects might be more than offset by the contract year end effects.

Overall, we find that bunching generally persists after taking into account the impact of potential reference points. Hence we believe bunching is mainly driven by the retention effect of options, rather than solely by reference points.

#### **4.3.4 Falsification test: absence of bunching for involuntary exits**

We conclude this section with a falsification test using option owners who leave the firm involuntarily. Involuntary exits can be the result of firing or layoffs, and leave the employee with barely any control over the timing of her exit. If bunching is driven by confounding factors that are orthogonal to option vesting, then it will likely show up among involuntary exits as well. If bunching is indeed due to the retention effect of options, then it would not be present in the distribution of involuntary exits.

Figure 9 shows involuntary exits of option owners over time. Unlike voluntary quitting, the distribution of involuntary exits shows no discernible bunching around the vesting dates.

[FIGURE 9 ABOUT HERE]

We then formally estimate the bunching model and illustrate the estimates in Figure 10. We set the affected ranges to be the same as the baseline estimates to accommodate the smaller sample of involuntary exits. The estimated counterfactual distributions fit the observed distribution fairly well. But once again we find no bunching at the vesting dates. The estimated excess mass is either not significantly different from 0 or even negative. The absence of bunching in the distribution of involuntary exits rejects the hypothesis that the baseline bunching effects are driven by confounding factors.

[FIGURE 10 ABOUT HERE]



## 5 Extension: Early Quitters and Benefit-Cost Analyses

In this section, we extend the baseline results to discuss two issues that are outside of the bunching framework, yet relevant to understanding employee stock options. We first look at option owners who voluntarily quit shortly before vesting dates. We explore the possibility of optimization error, which could restrict the retention benefit of options. We then discuss whether the retention effect is economically important for the firm. We answer this question in two steps. We first evaluate the quality of retained option owners by looking into the performance evaluation scores. Superior performance of retained option owners is a prerequisite for assessing any retention *benefit* of options. Then we do a back-of-the-envelope analysis of the benefit and cost of options.

### 5.1 Quitting before vesting

As discussed in Section 3, postponing quits when a vesting date is drawing near can greatly reduce the cost of quitting. But there are a small but non-negligible group of employees who chose to quit shortly *before* vesting dates. Because all unvested shares revert to the firm, these early quitters are essentially leaving money on the table—large amounts of money if the option values are high. We study the characteristics of early quitters and explore what drives the curious timing of their exits. In particular, we are interested in the possibility of optimization frictions.

We focus on the sample of option owners who quit within the affected ranges, both before and after the vesting dates. And for the empirical tests that follow, we define *early quitters* as those who quit before options vest, i.e.  $z_-^*$  months before the vesting date.<sup>14</sup> We pay special attention to a subgroup of early quitters, those who quit within one month preceding a vesting date. We compare the characteristics of the early quitters with those who quit after the same vesting date but still within the affected range.

[TABLE 5 ABOUT HERE]

---

<sup>14</sup>Recall that  $z_-^*$  is the baseline estimate of the number of months in affected range that are before the vesting date.

Table 5 compares the characteristics of the three groups of quitters. The early quitters hold options of much lower value than their peers who delayed quitting. The relative value of options is 22-23 times the early quitters' annual base wage, but is over 70 times for the delayed quitters. Although early quitters do have about 50% higher wage, the absolute value of their options is still substantially lower. The low relative and absolute values of options may explain why the owners do not wait for the upcoming vesting dates.

The early quitters also hold modestly higher level positions when they quit, with an average level of 6, or senior department managers, The average delayed quitter is at level 5, or junior department managers. We then compare the work experience of early quitters with that of delayed quitters, because the higher positions could be a result of longer tenure within the firm. And early quitters may be reluctant to delay quitting because they have already worked in the firm for an extended period of time. We find this not to be the case: early quitters have slightly more general work experience, but slightly less experience within this firm.

Finally, the early quitters do not differ from their peers who delayed quitting in gender, age, or education attainment. Hence it is unlikely that early quitting is driven by a lack of knowledge or inability to fully optimize.

In light of the comparison results, a plausible rationalization of early quitting is better outside offers. Early quitters have higher-level positions and higher wages, which reflect stronger abilities that are also valued by other firms. So they may get attractive outside offers that could make up for the foregone option values. One caveat about the comparison in Table 5, though, is that the estimates are rather noisy due to the very small sample size. Nonetheless, it could be a helpful first step in exploring reasons why some owners quit right before the options vest.

## **5.2 Quick benefit-cost analysis**

### **5.2.1 Performance of option owners**

Before assessing the retention "benefits" of options with a vesting schedule, we first look into the performance of retained option owners. Keeping employees in the firm does

not generate value in itself. So we need to verify whether the retained option owners are indeed beneficial to the firm.

We compare the performance of option owners between 2004 and 2012 in Figure 11. All employees receive a performance score at the end of each quarter, which is a number between 0 and 1 assigned by the department director.<sup>15</sup> The scores do not translate into a concrete measure of performance such as revenues or attendance, and should be interpreted in relative terms only.

[FIGURE 11 ABOUT HERE]

We first examine option owners who are retained within our sample period, and compare them with those who leave. The two top panels plot the results. Performance scores of option owners who do not quit remain steady over time for Plan 1 owners, and even increase slightly for Plan 2 owners. They have similar performance with owners who quit voluntarily, and both groups dominate option owners who leave involuntarily.

Then we compare the performance of option owners and non-owners. The bottom-left panel shows that option owners who stay and who quit both outperform non-owners. The superior performance of option owners is consistent with the above-noted endogeneity of option ownership.

The comparisons in Figure 11 show that the retained option owners are in general more valuable to the company, whether they are retained for the mere short term or for longer. Hence we have established the premise for assessing the retention benefit of options. We now proceed to the benefit-cost analysis.

### 5.2.2 Benefit-cost analysis

We do a simple, back-of-the-envelope benefit-cost analysis to see whether the retention effect of options is enough to offset the cost of options to the company. We first calculate the Black-Scholes value of options granted to each employee  $i$ , and use it to measure the firm's granting cost, denoted by  $c_i$ . We then calculate the benefit of options, which comes

---

<sup>15</sup>Employees who receive piece-rate pays do not have performance scores, and are thus not included in the sample.

from the following two ways.

The first benefit of granting options is they partially substitute wages and help the firm save some wage expenses. Without options, the firm will need to pay higher wages to the option owners. Estimating the counterfactual wage, however, is challenging because of the endogeneity problem. Hence a precise estimate is out of the scope of this paper. For a quick benefit-cost analysis, we use simple extrapolation to get a rough estimate: we first calculate the average monthly wage growth rate of non-owners,  $\bar{g}_t$ ; then we extrapolate option owner  $i$ 's counterfactual wage  $\hat{w}_{it}$  as

$$\hat{w}_{it} = \hat{w}_{i,t-1}(1 + \bar{g}_{t-1}), \quad t = 1, 2, \dots, T \quad (6)$$

where  $\hat{w}_{i0}$  is the same as  $w_{i0}$ , the observed wage of  $i$  in the month of granting.  $T$  is the month of  $i$ 's exit or month 60, whichever comes earlier.

We then take the difference between the observed monthly wage,  $w_{it}$ , and the counterfactual monthly wage,  $\hat{w}_{it}$ , and sum the differences from the granting month to month  $T$ . That is

$$S_i = \sum_{t=1}^T (\hat{w}_{it} - w_{it}) \quad (7)$$

which is the estimate of the firm's wage savings from granting options to employee  $i$ .

The second benefit of options comes from the increased output by the retained option owners who have superior performance. Because the performance scores do not directly measure employee output, we need to convert them into pecuniary values first. We start by regressing non-owner wages on their performance scores. Then we take the estimated coefficient as a conversion factor, and translate the performance score of option owners into a "fair wage,"  $\hat{y}_{it}$ . In doing so, we are implicitly assuming that wages are set to pay for an employee's output. It may be counterfactual, but it helps us quantify the value of an employee's performance to the firm. Finally, we take the difference between the fair wage  $\hat{y}_{it}$  of option-owners and the average wage of non-owners,  $\bar{y}_t$ . We sum the difference to get the total retention benefit,  $B_i$ .

One complication in calculating the retention benefit,  $B_i$ , is to identify over which periods the wage differences should be summed. We first propose a lower-bound measure,  $B_i^L$ :

$$B_i^L = \sum_{t=0}^{T_i} (\hat{y}_{it} - \bar{y}_t) + B^* \times \sum_{t=z^*-z_-^*}^{z^*+z_+^*} \bar{y}_{it} \quad (8)$$

The lower-bound retention benefit measure consists of two parts. First, it includes the sum of wage differences ( $\hat{y}_{it} - \bar{y}_t$ ) between month 0 and month  $T_i$ , the final period of  $i$  in the sample. This term captures the value of increased outputs by option owners. Second,  $B_i^L$  also takes into account the sum of fair wages ( $\bar{y}_{it}$ ) over the affected range around the *most recent* vesting date. It captures the benefits of keeping the option owner, who could have quit in these periods. The sum is then adjusted by the size of bunching,  $B^*$ , which is the excess quitting after the vesting dates. The adjustment reflects the fact that not all option owners necessarily quit after vesting dates. For those who do not quit, the extensive margin benefit is zero and should not be included in  $B_i$ . The more bunching there is, the more excess quitting after option vesting, and the more likely  $i$  will quit after the vesting date, in which case  $i$ 's output needs to be included in the extensive margin benefits. The first component in (8) is the *intensive* margin benefit from higher outputs created by option owners relative to non-owners. The second component is the *extensive* margin benefit from having any output created by option owners relative to losing the option owners.

We also propose an upper-bound measure of the retention benefit,  $B_i^U$ :

$$B_i^U = \sum_{t=0}^{T_i} (\hat{y}_{it} - \bar{y}_t) + \text{Retained} \times \text{Exercised} \times \left( \sum_{t=z^*}^{T_i} \bar{y}_{it} \right) + \text{Quit} \times \left( \sum_{t=z^*-z_-^*}^{T_i} \bar{y}_{it} \right) \quad (9)$$

the value of which depends on the option owner's turnover and option exercising. The first component is common to all option owners and is the same as in 8. The second component captures the additional benefit from option owners who are retained in the firm *and* have exercised at least 70% of vested shares since the last vesting date. We consider these owners to have the potential intention of quitting, thus their output (quantified by the fair wage)

from that point onward adds to the firm's benefit from keeping these employees. The last component is unique to option owners who quit. Because a quitting option owner might have quit any time in the absence of vesting dates, we consider the additional output during the quitting owner's extended stay as part of the firm's retention benefit. Hence the term sums the owner's fair wage since the beginning of the affected range until the actual time of quitting.

[TABLE 6 HERE]

We compare the median benefits and cost by option plan and employee job level in Table 6. For Plan 1 owners, the average cost of granting options is about 1,609 thousand yuan per person. The wage savings are 1153 thousand yuan, or about 2.8 times of the granting cost. The retention benefit is between 74 and 329 thousand yuan, which is another 10-94% of the granting cost. The two types of benefits combined are about 3.90-3.75 times the granting cost. Plan 1 owners in lower levels have a substantially higher benefit-cost ratio, which is mainly the artifact of high wage savings. The benefit-cost ratio is lower for level 5-10 owners, ranging from 0.42 to 1.06.

The benefit-cost comparison on Plan 2 owners is qualitatively similar. For Plan 2 owners as a whole, the benefit from wage savings alone is 1.9 times the granting cost. The retention benefit estimates have a wide range, with the lower bound being 7% of the cost and the upper bound 147%. The total benefits also outweigh the cost, as is the case with Plan 1 options. Overall benefit-cost ratio is between 1.95-3.34.

The benefit-cost comparison shows that the benefit of granting options is sufficient to justify the cost. In addition, the retention benefits contribute sizable shares of the total benefit. And the retention benefits alone can be large enough to justify the granting cost in some circumstances. Hence accounting for the retention effects of options can avoid the potential underestimation of the benefit-cost ratio, thereby helping to reconcile the popularity of options and the high granting costs.

## 6 Concluding Remarks

In this paper, we use the bunching design to identify the causal effect of employee stock options on turnover. We take advantage of the discontinuous changes in the cost of quitting around option vesting dates, and estimate the bunching model with a novel employee-level panel data. We find significant and robust bunching at the vesting dates, which quantifies the retention effect of option ownership. We also show that the retention effects are stronger for options with higher values, nonexistent for involuntary exits, and unlikely to be driven by confounding factors such as reference points.

The immediate policy-relevant question stemming from the findings is whether the benefit of options is large enough to justify their costs. To that end, we discuss the performance of option owners and find that the retained employees are indeed valuable to the company. We also do a simple benefit-cost analysis and show that the retention benefits contribute greatly to the total benefit of options, which more than justifies the granting cost on average.

One caveat for interpreting the estimates is the external validity. While the bunching design is known for its strong internal validity that identifies the treatment effect, the results may not be readily generalizable to other settings. Also, note that the retention effects captured by bunching are identified from employees who changed the *time to quit* because of options. But those who changed the *decision to quit* and stayed in the firm for good are observationally equivalent to those who never wanted to quit. Hence we could not separately identify the retention effect on option owners who give up quitting, and leave the question for future work.

Nonetheless, our findings shed light on the puzzling popularity of broad-based employee stock options, in spite of high granting costs. Significant retention benefits provide an additional reason for the prevalence of options, which the incentivization hypothesis alone is hard to explain. Bunching at the vesting dates also suggests why there is a vesting schedule in the first place. Given that options hardly induce long-term efforts, having a vesting schedule can at least retain option owners for the short term.

## References

- Aggarwal, R. K. and A. A. Samwick (1999). The Other Side of the Trade-off: The Impact of Risk on Executive Compensation. *Journal of Political Economy* 107(1), 65–105.
- Barron, J. M. and R. Waddell (2003). Hard Work Versus Good Intentions: Stock Options as Compensation. *Purdue University and University of Oregon Working Paper*.
- Carter, M. E. and L. J. Lynch (2004, February). The Effect of Stock Option Repricing on Employee Turnover. *Journal of Accounting and Economics* 37(1), 91–112.
- Core, J. E. and W. R. Guay (2001). Stock Option Plans for Non-executive Employees. *Journal of Financial Economics* 61, 253–287.
- Cowgill, B. and E. Zitzewitz (2014). Incentive Effects of Equity Compensation: Employee-level Evidence from Google. *Dartmouth Department of Economics Working Paper*.
- Dube, A. and R. B. Freeman (2010). Complementarity of Shared Compensation and Decision-Making Systems: Evidence from the American Labor Market. In D. Kruse, R. B. Freeman, and J. Blasi (Eds.), *Shared Capitalism at Work: Employee Ownership, Profit and Gain Sharing, and Broad-based Stock Options*, Chapter 5, pp. 167–199. University of Chicago Press.
- Hall, B. J. and K. J. Murphy (2003). The Trouble with Stock Options. *Journal of Economic Perspectives* 17(3), 49–70.
- Hausman, C. and D. S. Rapson (2017). Regression Discontinuity in Time: Considerations for Empirical Applications. *NBER Working Paper 23602*.
- Huddart, S. and M. Lang (1996, February). Employee Stock Option Exercises an Empirical Analysis. *Journal of Accounting and Economics* 21(1), 5–43.
- Imbens, G. and T. Lemieux (2008). Regression Discontinuity Designs: A Guide to Practice. *Journal of Econometrics* 142(2), 615–635.



- Kedia, S. and A. Mozumdar (2002). Performance Impact of Employee Stock Options. *Rutgers Business School Working Paper*.
- Kleven, H. J. (2016). Bunching. *Annual Review of Economics* 8, 435–464.
- Kleven, H. J. and M. Waseem (2013). Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan. *The Quarterly Journal of Economics*, 669–723.
- Kruse, D. L. (1993). *Profit Sharing: Does It Make A Difference?*, Chapter 1, pp. 1–21. Upjohn Institute Press.
- Lazear, E. P. (2000, February). Performance Pay and Productivity. *The American Economic Review* 90(5), 1346–1361.
- Lazear, E. P. (2004). Output-Based Pay: Incentives, Retention or Sorting? *Research in Labor Economics* 23, 1–25.
- Mehran, H. and J. Tracy (2001). The Impact of Employee Stock Options on the Evolution of Compensation in the 1990s. *Economic Policy Review* 3, 17–34.
- Oyer, P. (2004, August). Why Do Firms Use Incentives That Have No Incentive Effects? *The Journal of Finance* 59(4), 1619–1650.
- Oyer, P. and S. Schaefer (2005, April). Why Do Some Firms Give Stock Options to All Employees? An Empirical Examination of Alternative Theories. *Journal of Financial Economics* 76(1), 99–133.
- Oyer, P. and S. Schaefer (2006). Costs of Broad-Based Stock Option Plans. *Journal of Financial Intermediation* 15(4), 511 – 534.

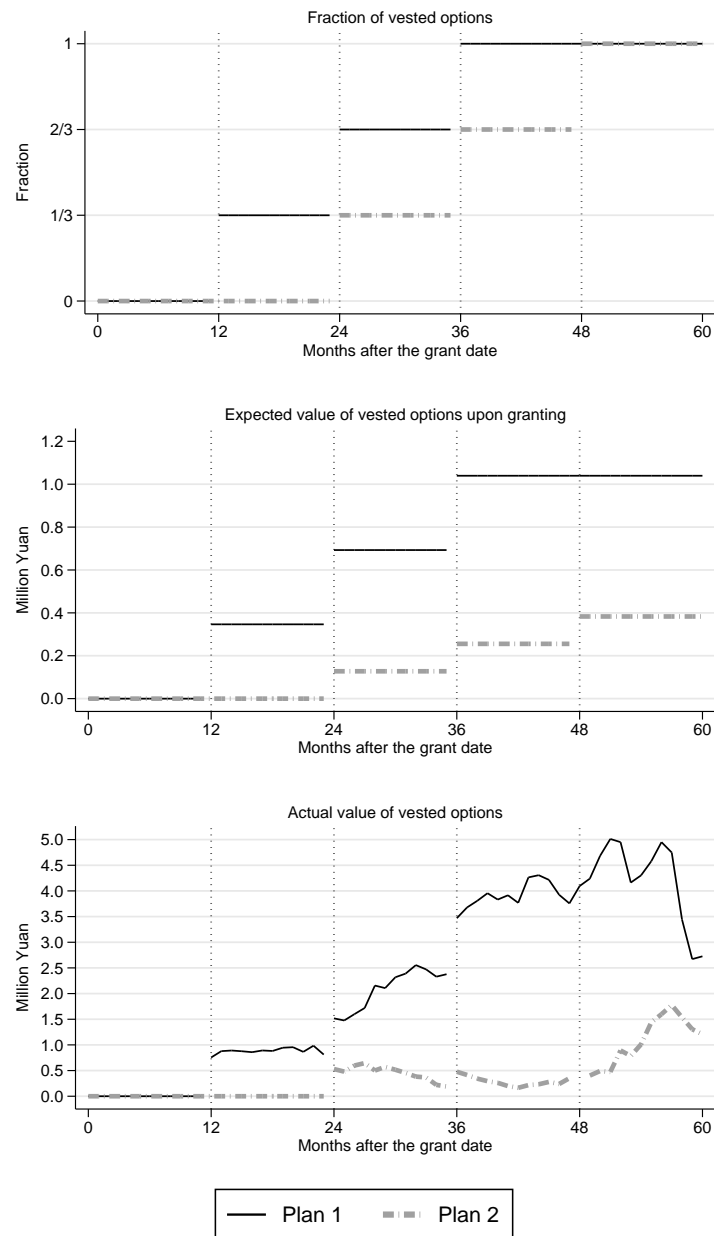


Figure 1: Option vesting schedules and values over time

NOTES: The horizontal axes show months after option granting. The grant dates are denoted as the beginning of month 0. Plan 1 options vest in three equal batches at the beginning of months 12, 24, and 36. Plan 2 options vest in three equal batches at the beginning of months 24, 36, and 48. Vesting dates are marked by dotted vertical lines. In all three panels, the solid (dash-dotted) line shows Plan 1 (Plan 2) averages. The top panel illustrates the vesting schedules by showing the fraction of vested options over time. The middle panel shows the expected value of vested shares, which is calculated by multiplying the number of vested shares at each point in time with the Black-Scholes value per share *on the grant date*. The bottom panel shows the actual value of vested shares by plotting their real-time Black-Scholes value in each month after granting.

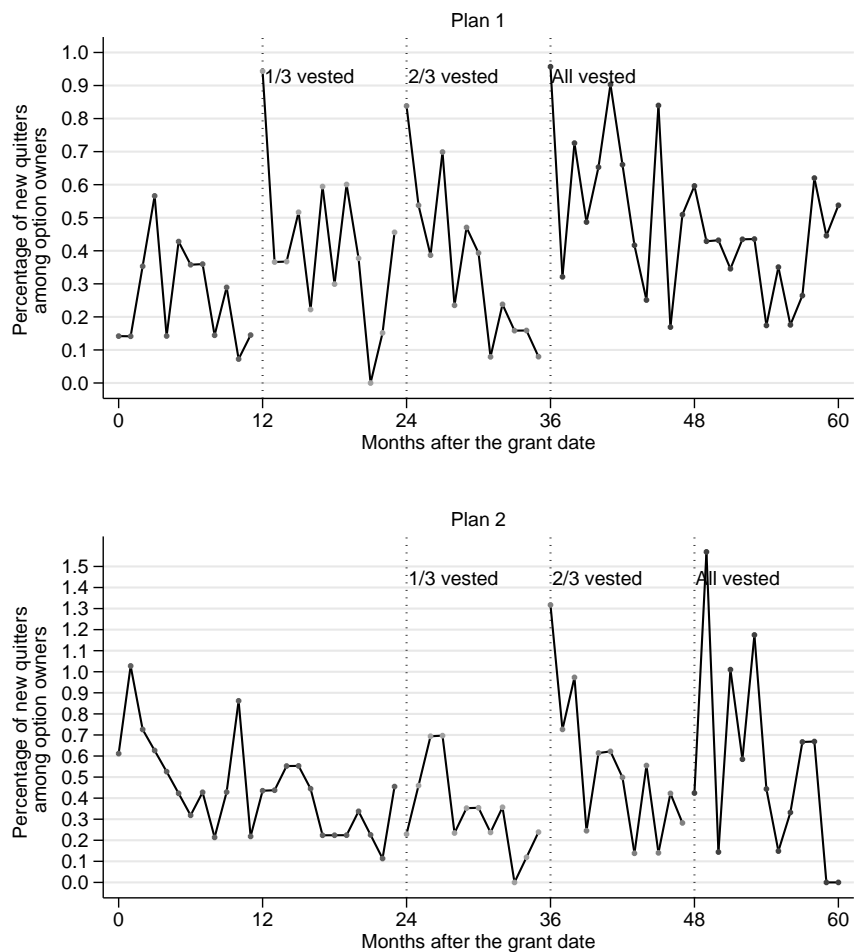


Figure 2: Preliminary evidence of bunching: monthly quitting rates of option owners

NOTES: The horizontal axes show months after option granting, with the grant date being the first day of month 0. Option vesting dates are marked by dotted vertical lines (months 12, 24, and 36 for Plan 1; months 24, 36, and 48 for Plan 2). The vertical axes show monthly quit rates (in %), calculated as the fraction of voluntary quitters in each month as a fraction of total initial option owners of that plan.

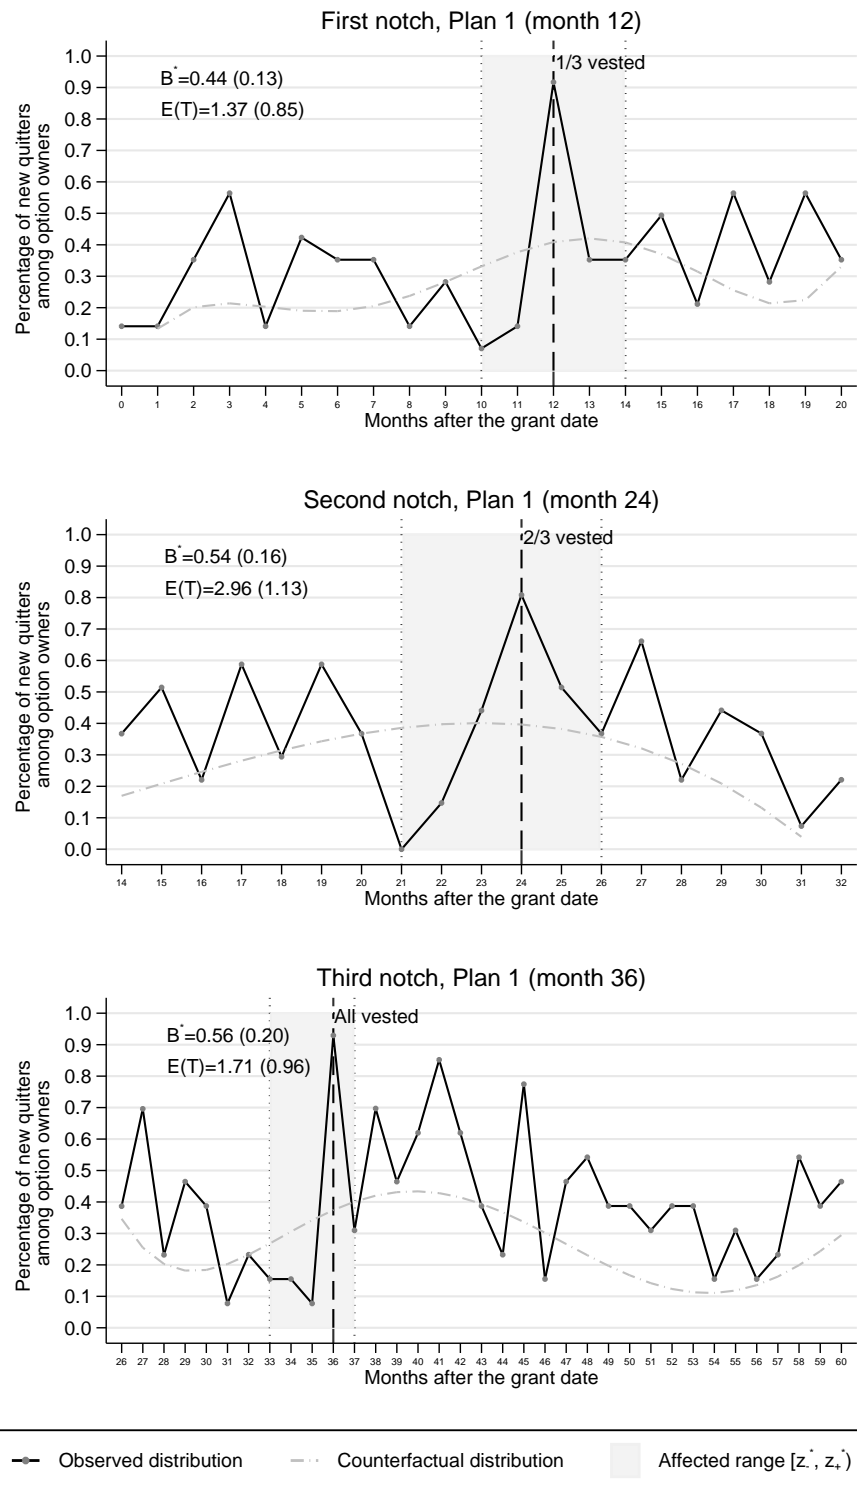


Figure 3: Baseline estimates of bunching: empirical and counterfactual distributions of quitters and the affected range (Plan 1)

NOTES: The horizontal axes show months after option granting, with the grant date being the first day of month 0. The three panels show bunching at the three vesting dates for Plan 1. The solid lines trace out the observed distribution of quitters. The dashed lines show the counterfactual distributions. The shaded area indicates the affected range. Sizes of bunching at the three vesting dates, measured by excess mass  $B^*$ , are reported at the top-left corner of each plot. The corresponding standard errors are in parentheses. Also reported are the expected number of months an option owner is retained,  $E(T)$ .

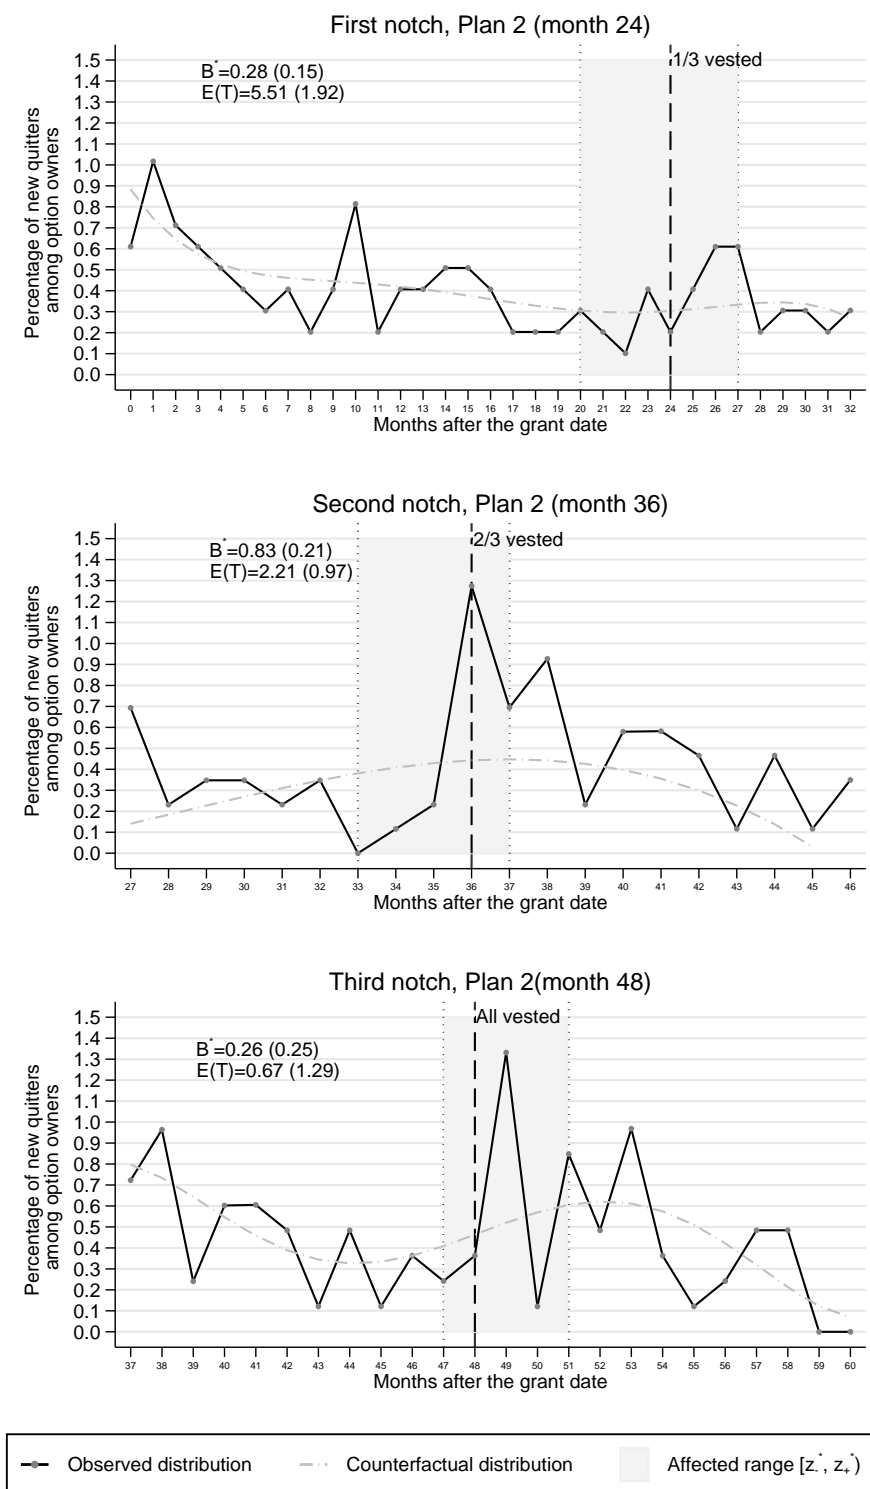


Figure 4: Baseline estimates of bunching: empirical and counterfactual distributions of quitters and the affected range (Plan 2)

NOTES: The horizontal axes show months after option granting, with the grant date being the first day of month 0. The three panels show bunching at the three vesting dates for Plan 2. The solid lines trace out the observed distribution of quitters. The dashed lines show the counterfactual distributions. The shaded area indicates the affected range. Sizes of bunching at the three vesting dates, measured by excess mass  $B^*$ , are reported at the top-left corner of each plot. The corresponding standard errors are in parentheses. Also reported are the expected number of months an option owner is retained,  $E(T)$ .


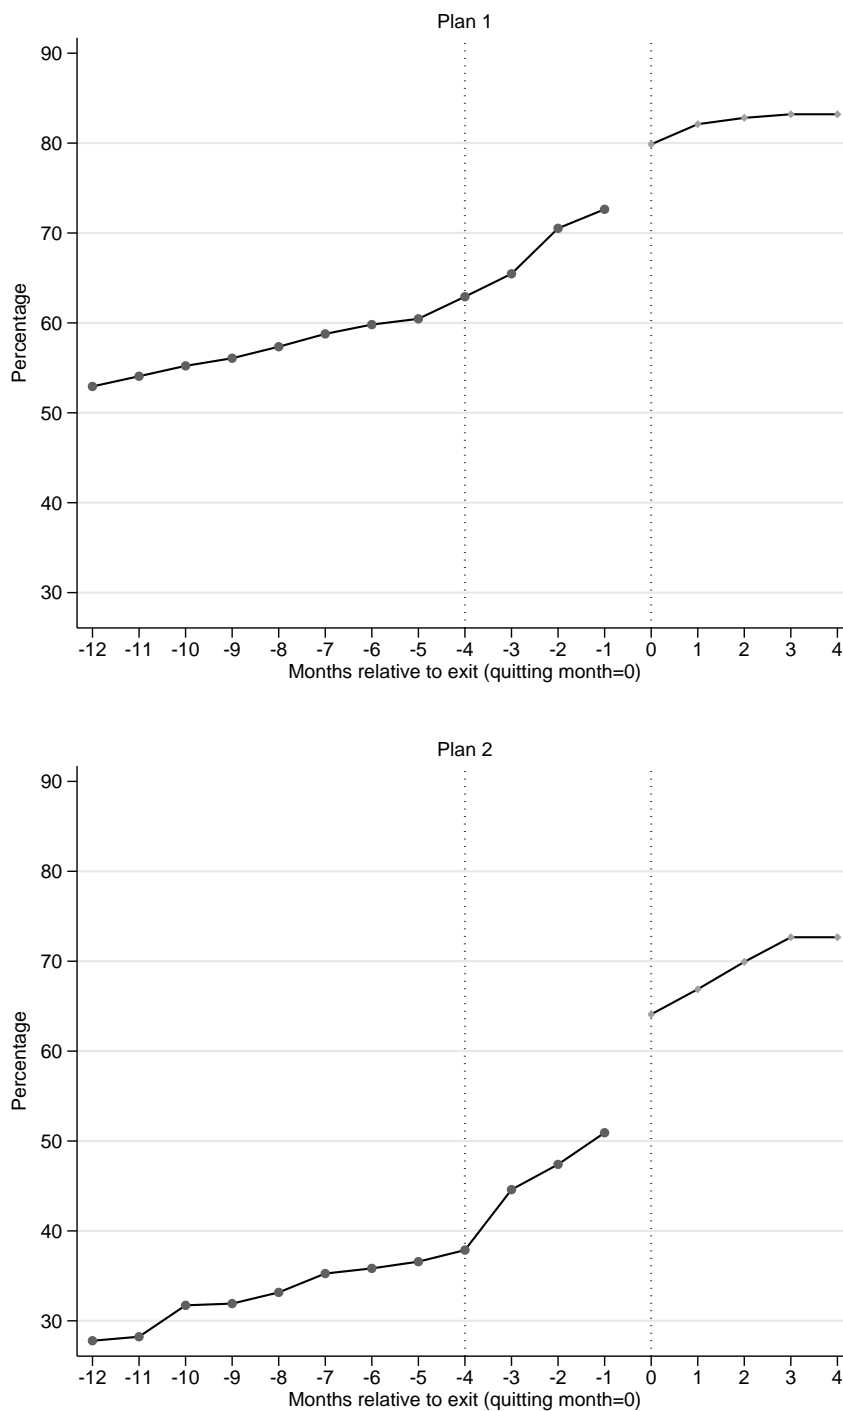


Figure 5: Percentage of option owners exercising all vested shares over time

NOTES: The solid lines plot the percentage of option owners exercising all vested options in each given month. It is the number of option owners who exercise at least the number of shares that have most recently vested, divided by the total number of option owners still in the firm in that month.



**Figure 6: Cumulative fraction of exercised options by quitting owners**

NOTES: The horizontal axes are months relative to an option owner's exit. Month 0 is the month of quitting; months with negative indices are those preceding the quit; months with positive indices are those after the quit. For option owners who quit, we track the cumulative number of options exercised as a fraction of total shares initially granted over time. The solid lines plot the fractions averaged over quitting owners. The fractions do not reach 100% because some option owners quit while still having unvested options, which revert to the firm immediately. The fractions remain unchanged after month 3 because quitting owners can exercise vested options only within 90 days after they quit.




Figure 7: Testing for reference points: distribution of quitting within contract and calendar years

NOTES: The top panel plots the distribution of quitting time within a contract year, with the ending month of contract years being month 0, and the  $s$ th month of contract years being month  $s$ . Each shaded (unshaded) bar shows the fraction of option owners (non-owners) quitting in a given month among all owners (non-owners) who quit. Similarly, the bottom panel is the distribution of quitting time within a calendar year.



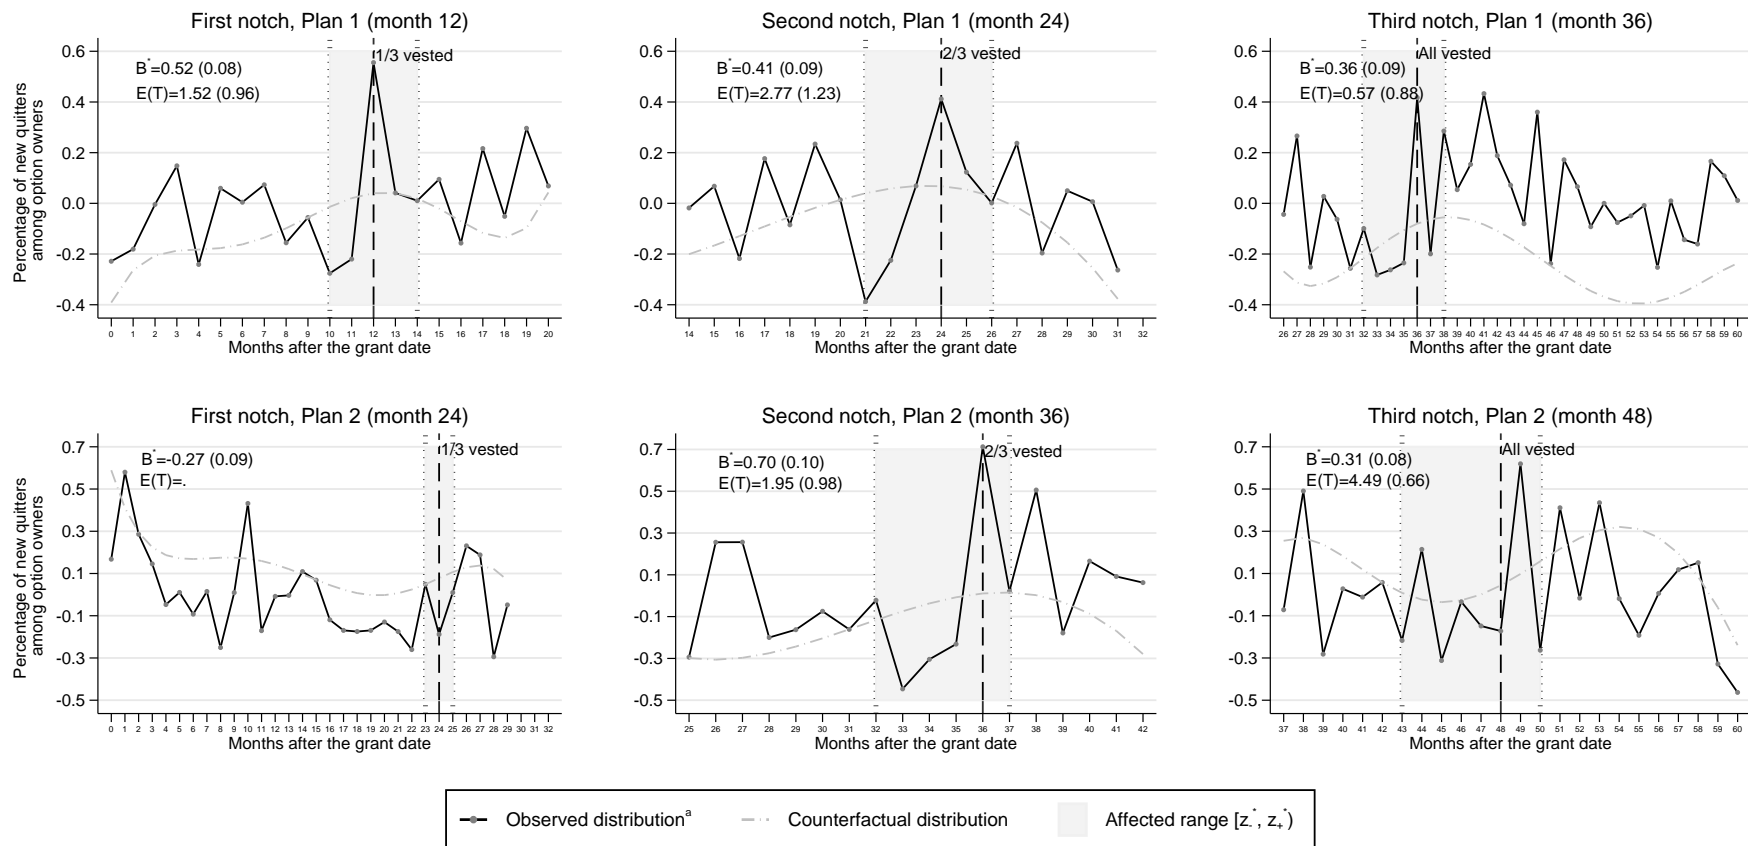


Figure 8: Testing for reference points: bunching estimates with month fixed effects

NOTES: <sup>a</sup>Observed distribution of quitters (solid line) is conditional on potential reference point effects (hence the negative values in some parts of the distribution). Potential reference points are February, May, June, September, and the first 3 months of a contract year. The horizontal axes show months after option granting, with the grant date being the first day of month 0. The dashed lines show the counterfactual distributions, estimated by fitting a flexible quintic polynomial. The shaded area indicates the affected range. The sizes of bunching, measured by excess mass  $B^*$ , are reported at the top-left corner of each plot. The corresponding standard errors are in parentheses. Also reported are the expected number of months an option owner is retained,  $E(T)$ .  $E(T)$  for the first notch of Plan 2 (month 24) is not well defined due to the absence of bunching.




Figure 9: Monthly involuntary exit rates of option owners

NOTES: The horizontal axis shows months after option granting, with the grant date being the first day of month 0. Vesting dates are marked by dotted vertical lines. The vertical axes show monthly rates of involuntary separations (in %). The solid lines plot the number of option owners who leave the firm involuntarily (e.g. fired or laid off) in each month as a fraction of all initial option owners.

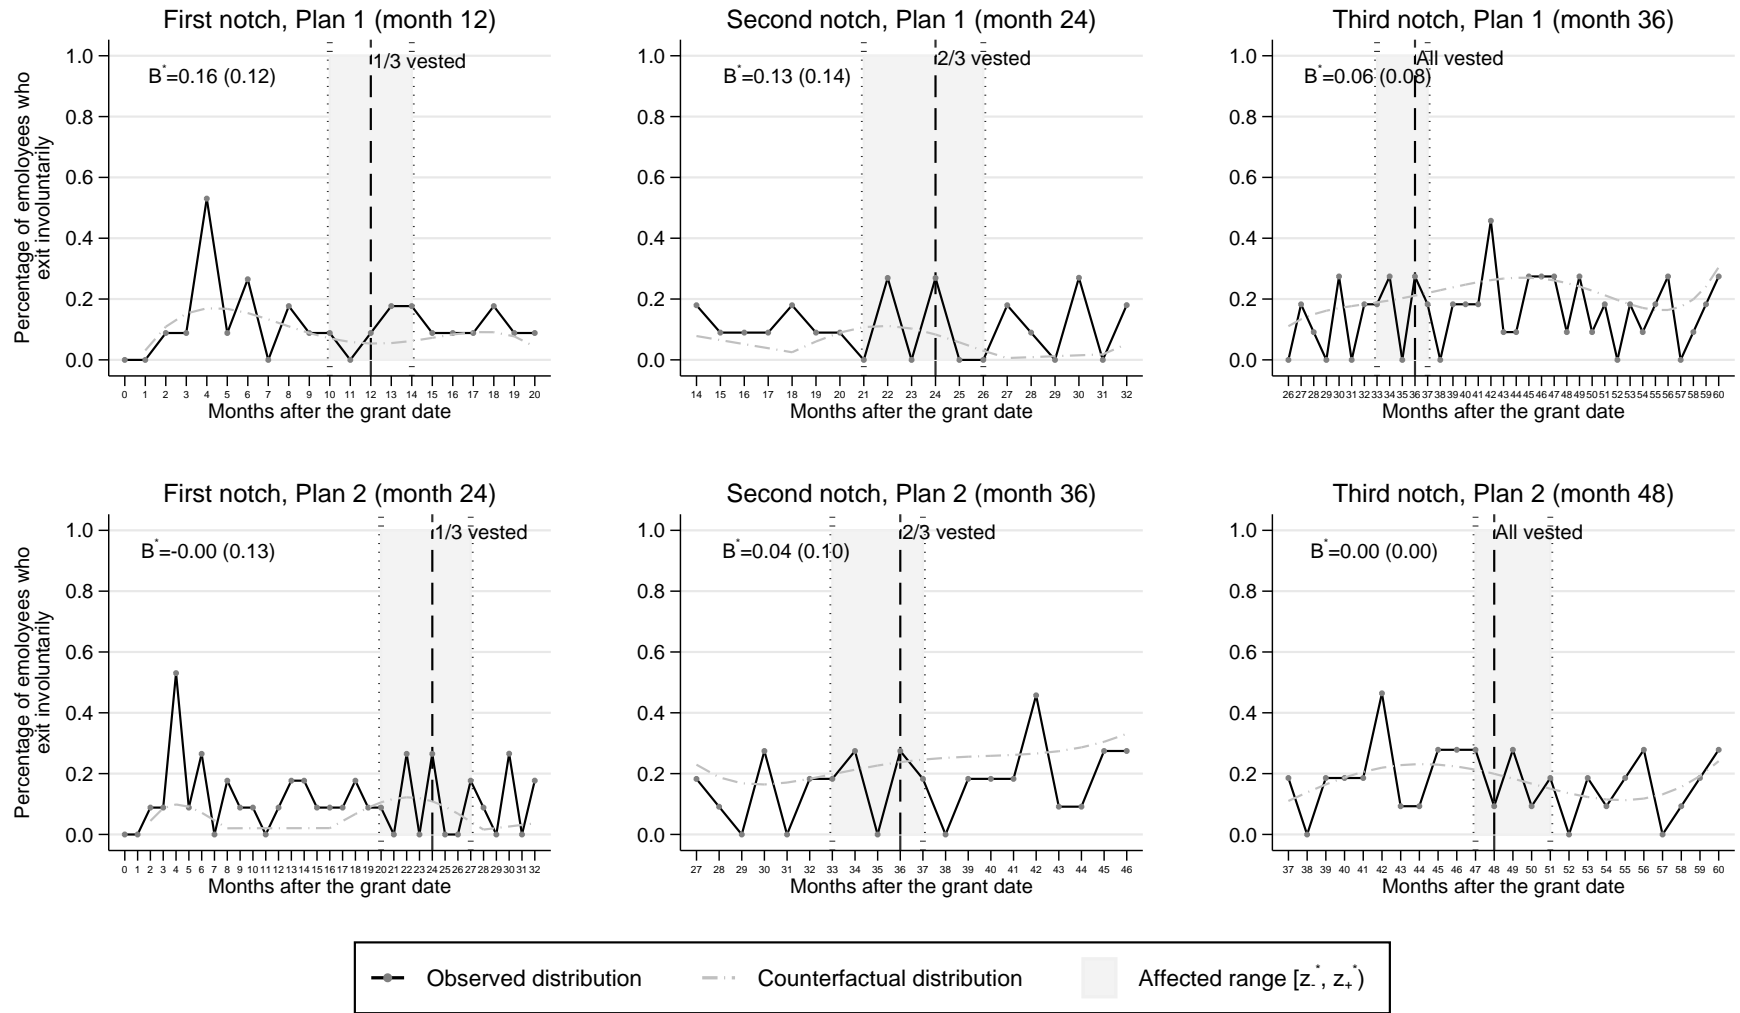


Figure 10: Falsification test: absence of bunching for involuntary exits

NOTES: The horizontal axes show months after option granting, with the grant date being the first day of month 0. The top, middle, and bottom plots show bunching at the first, second, and third vesting dates for Plan 1. The solid lines in the plots are the observed number of option owners who leave the firm involuntarily (e.g. fired or laid off) in each month as a fraction of all initial option owners. The dash-dotted lines are estimated counterfactual distributions of quitting rates over time in the absence of options, estimated by fitting a flexible quintic polynomial (using data points around the threshold, except those in the affected range). The shaded area indicates the affected range.

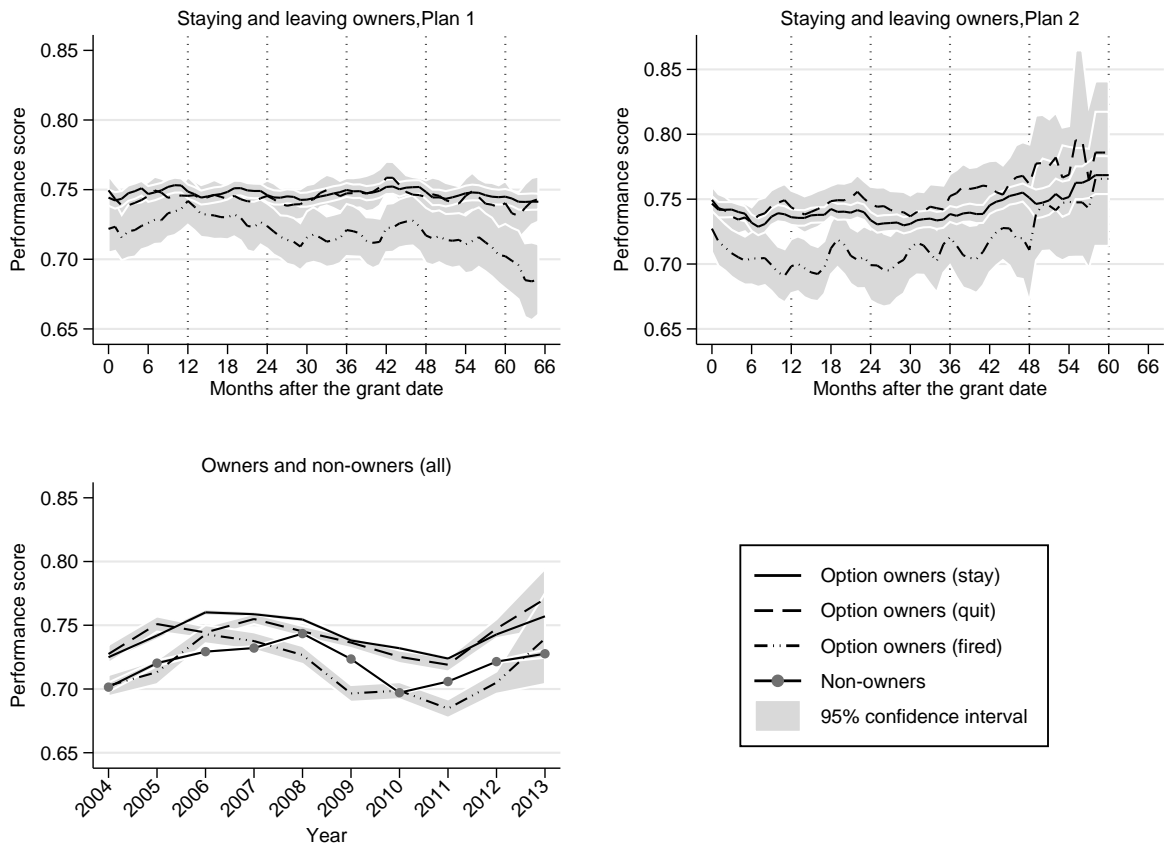


Figure 11: Performance evaluation scores of option owners

NOTES: We restrict our sample to salaried option owners from 2004 to 2012, during which time performance scores are available. Performance scores are numbers between 0 and 1 assigned quarterly by each employee's director. The horizontal axes in the two top panels are months after the granting date. These two panels plot the performance scores of option owners who stay (solid), quit (dashed), and leave involuntarily (dash-dotted). The bottom-left panel plots the performance scores of all employees over calendar time, including non-owners (solid with round marker) and the 3 types of options owners. The 95% confidence intervals are in shaded in gray for all plots.

	Plan 1				Plan 2			
	Mean	Std.Dev.	Min	Max	Mean	Std.Dev.	Min	Max
Shares granted (in 1,000)	9.124	24.527	0.532	672	6.394	9.982	0	128
(Option value)/(annual base wage)	14.550	14.860	0	223.902	5.373	7.067	0	111.138
Monthly base wage (in 1,000 yuan)	8.467	6.123	0.621	57.400	10.884	8.036	0.733	64.802
1{Hold other options}	0.638	0.481	0	1	0.660	0.474	0	1
1{Hold other Plan 1 options}	0.636	0.481	0	1	0.563	0.496	0	1
1{Hold other Plan 2 options}	0.101	0.302	0	1	0.390	0.488	0	1
1{Level 1}	0.048	0.214	0	1	0.052	0.222	0	1
1{Level 2}	0.048	0.214	0	1	0.065	0.246	0	1
1{Level 3}	0.158	0.365	0	1	0.053	0.224	0	1
1{Level 4}	0.190	0.392	0	1	0.048	0.213	0	1
1{Level 5}	0.230	0.421	0	1	0.350	0.477	0	1
1{Level 6}	0.192	0.394	0	1	0.254	0.435	0	1
1{Level 7}	0.075	0.263	0	1	0.100	0.300	0	1
1{Level 8}	0.039	0.194	0	1	0.044	0.205	0	1
1{Level 9}	0.017	0.131	0	1	0.026	0.159	0	1
1{Level 10}	0.002	0.050	0	1	0.009	0.094	0	1
Unique individuals	607				605			
Option-individual combinations	1,609				1,112			

Table 1: Option value and owner characteristics on grant dates

NOTES: All statistics are measured on the grant date. Option values are the Black-Scholes value, originally in US dollars and converted to real 2014 RMB yuan using the then-prevailing exchange rate and the province-level urban CPI in China.

Threshold month (1)	Bunching (excess mass $B^*$ ) (2)	Months retained ( $\mathbb{E}(T)$ ) (3)	Affected range		Sample size		
			Before ( $z_-^*$ ) (4)	After ( $z_+^*$ ) (5)	Observations (6)	Employees (7)	Option-employees (8)
<i>Panel A. Plan 1</i>							
12	0.44 (0.13)	1.37 (0.85)	2 (1.56)	2 (0.83)	28,949	529	1,426
24	0.54 (0.16)	2.96 (1.13)	3 (1.47)	2 (0.88)	24,912	507	1,365
36	0.56 (0.20)	1.71 (0.96)	3 (0.59)	1 (0.20)	42,171	473	1,299
<i>Panel B. Plan 2</i>							
24	0.28 (0.15)	5.51 (1.92)	4 (1.57)	3 (0.85)	29,961	529	985
36	0.83 (0.21)	2.21 (0.97)	3 (1.53)	1 (0.86)	16,283	449	873
48	0.26 (0.25)	0.67 (1.29)	1 (1.64)	3 (0.84)	17,052	431	838

Table 2: Joint estimates of bunching and the affected range: baseline specification

NOTES: Column (1) shows the threshold months (those with a vesting date) for which the results in other columns are estimated. Column (2) quantifies the magnitude of bunching, measured by the excess mass in the quitting distribution to the right of the threshold. Column (3) shows the expected length of retention due to option vesting in the threshold month. Columns (4)-(5) report the estimated affected range.  $z_-^*$  and  $z_+^*$  is the number of affected months before and after the vesting date, respectively. Columns (6)-(8) report the sample size used in the estimation for each threshold: the number of observations, the number of unique employees, and the number of unique option-employee combinations. Bootstrapped standard errors are in parentheses.

Threshold (1)	Black-Scholes value before vesting dates		Black-Scholes value when granted	
	Below median (2)	Above median (3)	Below median (4)	Above median (5)
<i>Panel A. Plan 1</i>				
12	0.728 (0.086)	0.721 (0.087)	0.536 (0.111)	0.494 (0.072)
24	0.075 (0.108)	0.567 (0.068)	0.308 (0.065)	0.800 (0.116)
36	0.594 (0.052)	0.758 (0.073)	0.637 (0.056)	0.695 (0.141)
<i>Panel B. Plan 2</i>				
24	-0.009 (0.039)	0.199 (0.024)	0.108 (0.104)	0.220 (0.062)
36	1.332 (0.295)	1.439 (0.265)	1.174 (0.107)	0.568 (0.230)
48	0.240 (0.135)	0.270 (0.156)	0.481 (0.150)	0.200 (0.377)

**Table 3: Heterogeneous bunching effects by option value**

NOTES: Column (1) shows the threshold months (those with a vesting date) for which the results in other columns are estimated. Columns (2)-(3) are using subsamples whose Black-Scholes value is below and above the median *before vesting dates*, respectively. Columns (4)-(5) are using subsamples whose Black-Scholes value is below and above the median *when granted*, respectively. Both measures of option value are calculated at the time of granting. The coefficients reported are estimates of the bunching effect, the excess mass  $B^*$ . Bootstrapped standard errors are in parentheses.

Threshold month (1)	Bunching (excess mass $B^*$ ) (2)	Months retained ( $\mathbb{E}(T)$ ) (3)	Affected range		Sample size		
			Before ( $z_-^*$ ) (4)	After ( $z_+^*$ ) (5)	Observations (6)	Employees (7)	Option-employees (8)
<i>Panel A. Plan 1</i>							
12	0.52 (0.08)	1.52 (0.96)	2 (1.42)	2 (0.80)	28,949	529	1,426
24	0.41 (0.09)	2.77 (1.23)	3 (1.87)	2 (0.82)	24,912	507	1,365
36	0.36 (0.09)	0.57 (0.88)	4 (0.71)	2 (0.21)	42,171	473	1,299
<i>Panel B. Plan 2</i>							
24	-0.27 (0.09)	- -	1 (1.14)	1 (0.60)	29,961	529	985
36	0.70 (0.10)	1.95 (0.98)	4 (1.38)	1 (0.79)	16,283	449	873
48	0.31 (0.08)	4.49 (0.66)	5 (0.91)	2 (0.51)	17,052	431	838

Table 4: Testing for reference points: bunching estimates with month fixed effects

NOTES: Column (1) shows the threshold months (those with a vesting date) for which the results in other columns are estimated. Column (2) quantifies the magnitude of bunching, measured by the excess mass in the quitting distribution to the right of the threshold. Column (3) shows the expected length of retention due to option vesting in the threshold month. Columns (4)-(5) report the estimated affected range.  $z_-^*$  and  $z_+^*$  is the number of affected months before and after the vesting date, respectively. Columns (6)-(8) report the sample size used in the estimation for each threshold: the number of observations, the number of unique employees, and the number of unique option-employee combinations. The econometric specification is the same as the baseline model, except we add fixed effects for February, May, June, and September, which are potential reference points in the calendar year; we also add fixed effects for the first, second and third month after the end of contract years to avoid underestimating bunching, because part of the retention effects may be offset by the reference point effects *in the absence of options* in these months. Bootstrapped standard errors are in parentheses.



	(1)	(2)	(3)	(4)
	Early quitters	Early quitters	Delayed quitters	Differences
	$[z^* - 1, z^*)$	$[z^* - z_-, z^*)$	$[z^*, z^* + z_+)$	(3)-(2)
(Option value)/(annual base wage)	23.226 (7.956)	22.441 (5.837)	70.134 (31.790)	47.694 (60.675)
Monthly base wage (yuan)	12044.631 (1404.108)	12158.156 (1296.772)	7780.083 (516.198)	-4378.073 (1183.422)
Job level	5.875 (0.473)	6.115 (0.365)	5.116 (0.202)	-0.999 (0.419)
Experience	14.205 (0.979)	13.155 (0.902)	12.679 (0.671)	-0.475 (1.296)
In-firm experience	4.618 (0.631)	4.806 (0.540)	5.494 (0.334)	0.689 (0.668)
Male	0.625 (0.125)	0.615 (0.097)	0.564 (0.057)	-0.051 (0.113)
Age	35.606 (1.025)	34.353 (0.881)	33.989 (0.598)	-0.365 (1.145)
Years of schooling	15.769 (0.231)	15.810 (0.148)	15.667 (0.192)	-0.143 (0.352)
Observations	17	31	86	

Table 5: Testing for optimization friction: comparison of early quitters and other quitters

NOTES: Column (1) shows statistics of option owners who quit in the month preceding a vesting date, i.e. between month  $(z^* - 1)$  and  $z^*$ . Column (2) shows statistics of option owners who quit within the affected range and before the vesting date, i.e. between month  $(z^* - z_-)$  and  $z^*$ . Column (3) shows statistics of option owners who quit within the affected range but after the vesting date, i.e. those who are bunching in  $[z^*, z^* + z_+)$ . Column (4) reports the differences between (3) and (2). All standard errors are in parentheses. Option value is the Black-Scholes value of options in months before the vesting date and within the affected range. Option values are originally in US dollars and converted to real 2014 RMB yuan using the then-prevailing exchange rate. All money variables are adjusted for inflation using province-level urban CPI in China. Experience is estimated from employee age and education attainment. In-firm experience is the number of years an employee worked in the firm when she quits, which is readily available in the data.

Level	Granting cost (1)	Wage savings (2)	Retention Benefits		Benefits as a fraction of granting cost					
			Min. (3a)	Max. (3b)	Wage savings (2)/(1)	Retention Benefits		Total benefits		
						Min. (3a)/(1)	Max. (3b)/(1)	Min. [(2)+(3a)]/(1)	Max. [(2)+(3b)]/(1)	
<i>Panel A. Plan 1</i>										
All levels	1609.17	1152.51	74.38	329.07	2.81	0.10	0.94	2.90	3.75	
Levels 1-4	733.26	2193.53	74.19	290.53	6.04	0.03	1.15	6.08	7.19	
Levels 5-10	2297.48	318.64	74.52	359.93	0.28	0.14	0.78	0.42	1.06	
<i>Panel B. Plan 2</i>										
All levels	811.99	322.58	27.93	284.39	1.87	0.07	1.47	1.95	3.34	
Levels 1-4	125.88	618.78	-19.54	172.38	7.50	-0.31	2.45	7.19	9.96	
Levels 5-10	996.61	242.88	40.70	314.52	0.35	0.18	1.20	0.53	1.55	

**Table 6: A simple benefit-cost analysis**

NOTES: All cost, savings, and benefit values are in thousand RMB yuan, and are adjusted for inflation using province-level urban CPI in China. The leftmost column shows the sample selection criteria: option owners of all levels; those in levels 1-4; those in levels 5-10. Column (1) shows the average cost of granting options to an individual employee in each subsample. Column (2) is the average estimated wage savings of the firm. Columns (3a) and (3b) report the minimum and maximum estimated retention benefit to the firm, respectively. The next columns present the various benefits as ratios to the granting cost: wage savings, minimum and maximum retention benefits, and minimum and maximum total benefits, which is the sum of wage savings and retention benefits.

## A Supplemental Summary Statistics

Table A1 reports summary statistics on the turnover pattern of option owners in the firm. Panel A shows the overall fraction of option owners who quit during the sample period. One-third of all Plan 1 owners quit, and the figures are very similar between lower-level owners and higher-level ones. The turnover rate is lower for Plan 2, with the average being 26.40%. The difference between plans is largely because Plan 1 options were granted several years before Plan 2 options. The turnover rates of Plan 2 owners also vary considerably by job level: lower-level employees tend to quit less frequently, with a turnover rate of 18.75%, whereas higher-level employees have a turnover rate of 28.44%.

Panel A. Turnover rates (overall and by level)				
	Plan 1		Plan 2	
	N	Quit%	N	Quit%
All	1426	33.31	985	26.40
Level 1-4	641	33.70	208	18.75
Level 5-10	783	32.82	777	28.44

Panel B. Survival function over time				
Years after granting	N	Remain%	N	Remain%
Year 1	1422	96.91	984	93.80
Year 2	1422	92.26	982	89.92
Year 3	1419	88.37	980	86.43
Year 4	1419	82.45	894	79.31
Year 5	1417	78.62	829	71.89
Year 6	1207	70.84		

Table A1: Descriptive statistics on turnover patterns

NOTES: The table summarizes the fraction of option owners who quit (Panel A) and the fraction of option owners who remain in the firm from 1 to 6 years after option granting.

Owners of the two types of option plans also quit at different rates over time. Plan 1 owners quit at a stable annual rate of about 4-6% for most years. Plan 2 owners quit at low rates in the first three years after option granting. But turnover speeds up from that point onwards.

## B Baseline Estimates with Larger Samples

We are aware of the limited sample size in our main analysis. And we take two different approaches to address this issue. Although both approaches have advantages and drawbacks, we show that the baseline finding remains: option owner turnover increase discontinuously once options vest.

Our first approach is to pool the three vesting points of each option, and estimate the corresponding bunching effects jointly. This increases the sample size, but at the same time increases the number of parameters to be estimated. We present the results in Figure A1.

Our second approach to addressing the small sample issue is to use weekly turnover rates. This significantly boosts the number of observations available for the estimation of bunching at each vesting dates. But the enlarged sample size comes at the expense of more noisy turnover measures, as *weekly* turnover rates are much more volatile than monthly ones. Figures A2 and A3 summarizes the results.

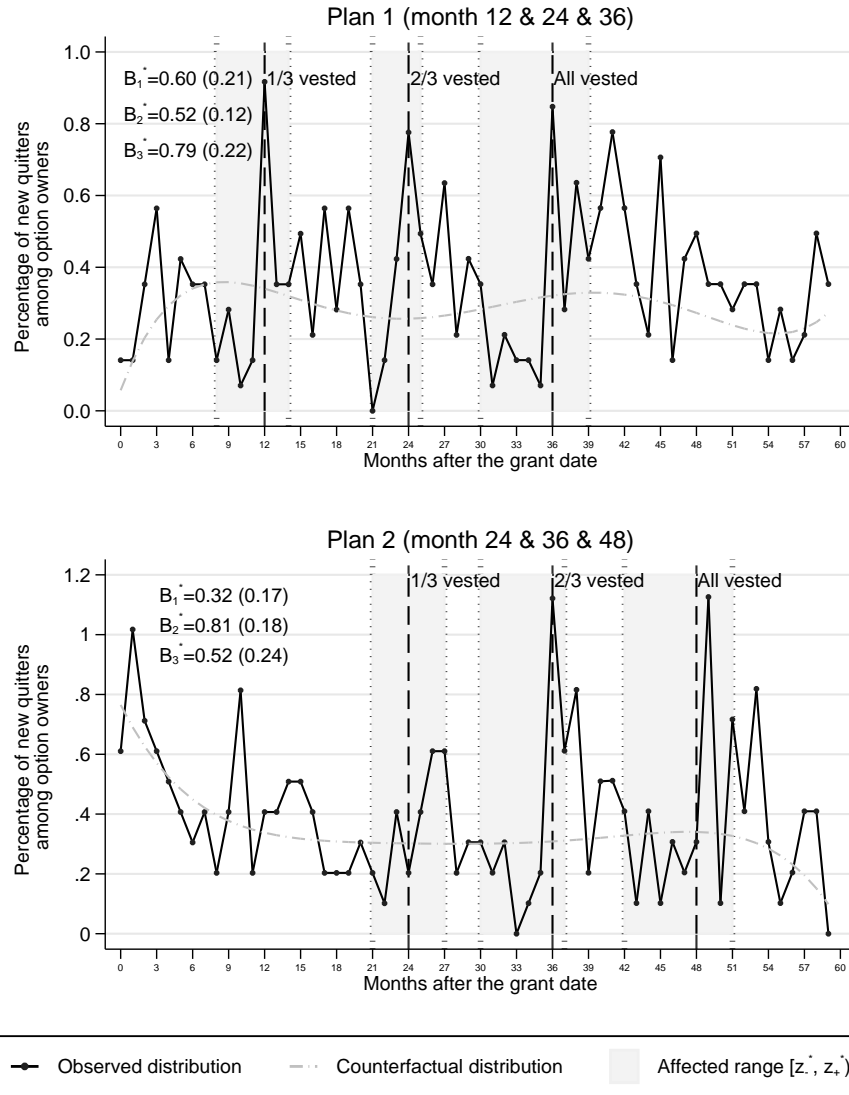


Figure A1: Pooled estimation of bunching at all three notches

NOTES: The horizontal axes show months after option granting, with the grant date being the first day of month 0. For each option plan, we assume all months after the grant date share the same counterfactual distribution. The solid lines trace out the observed distribution of quitters. The dashed lines show the counterfactual distributions. The shaded area indicates the affected range. Sizes of bunching at the three vesting dates, measured by excess mass  $B^*$ , are reported at the top-left corner of each plot. The corresponding standard errors are in parentheses. Also reported are the expected number of months an option owner is retained,  $E(T)$ .

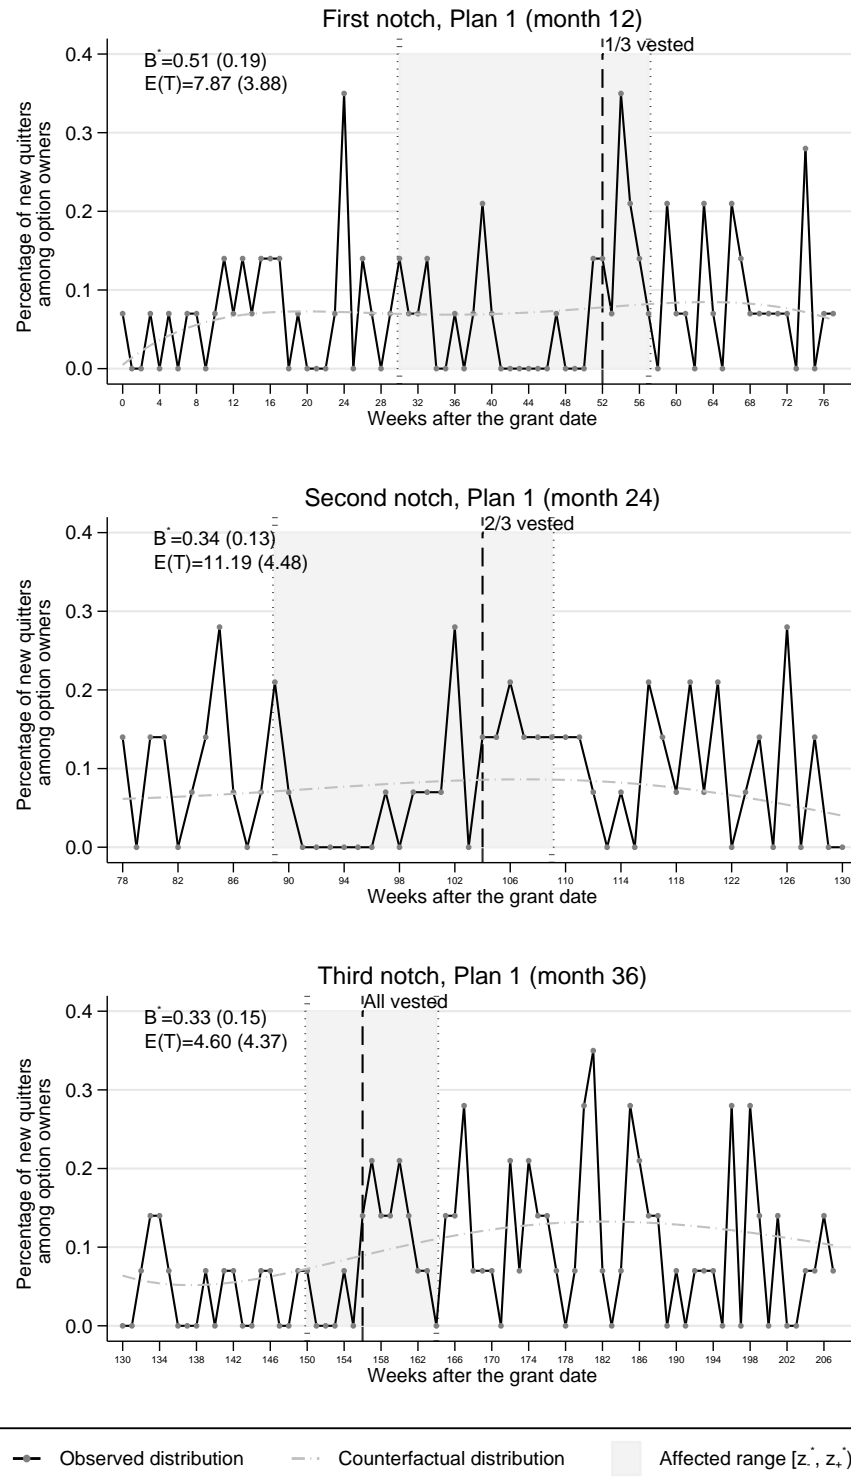


Figure A2: Estimation of bunching using weekly turnover rates (Plan 1)

NOTES: The horizontal axes show weeks after option granting, with the grant date being the first day of week 0. The three panels show bunching at the three vesting dates for Plan 1. The solid lines trace out the observed distribution of quitters. The dashed lines show the counterfactual distributions. The shaded area indicates the affected range. Sizes of bunching at the three vesting dates, measured by excess mass  $B^*$ , are reported at the top-left corner of each plot. The corresponding standard errors are in parentheses. Also reported are the expected number of months an option owner is retained,  $E(T)$ .

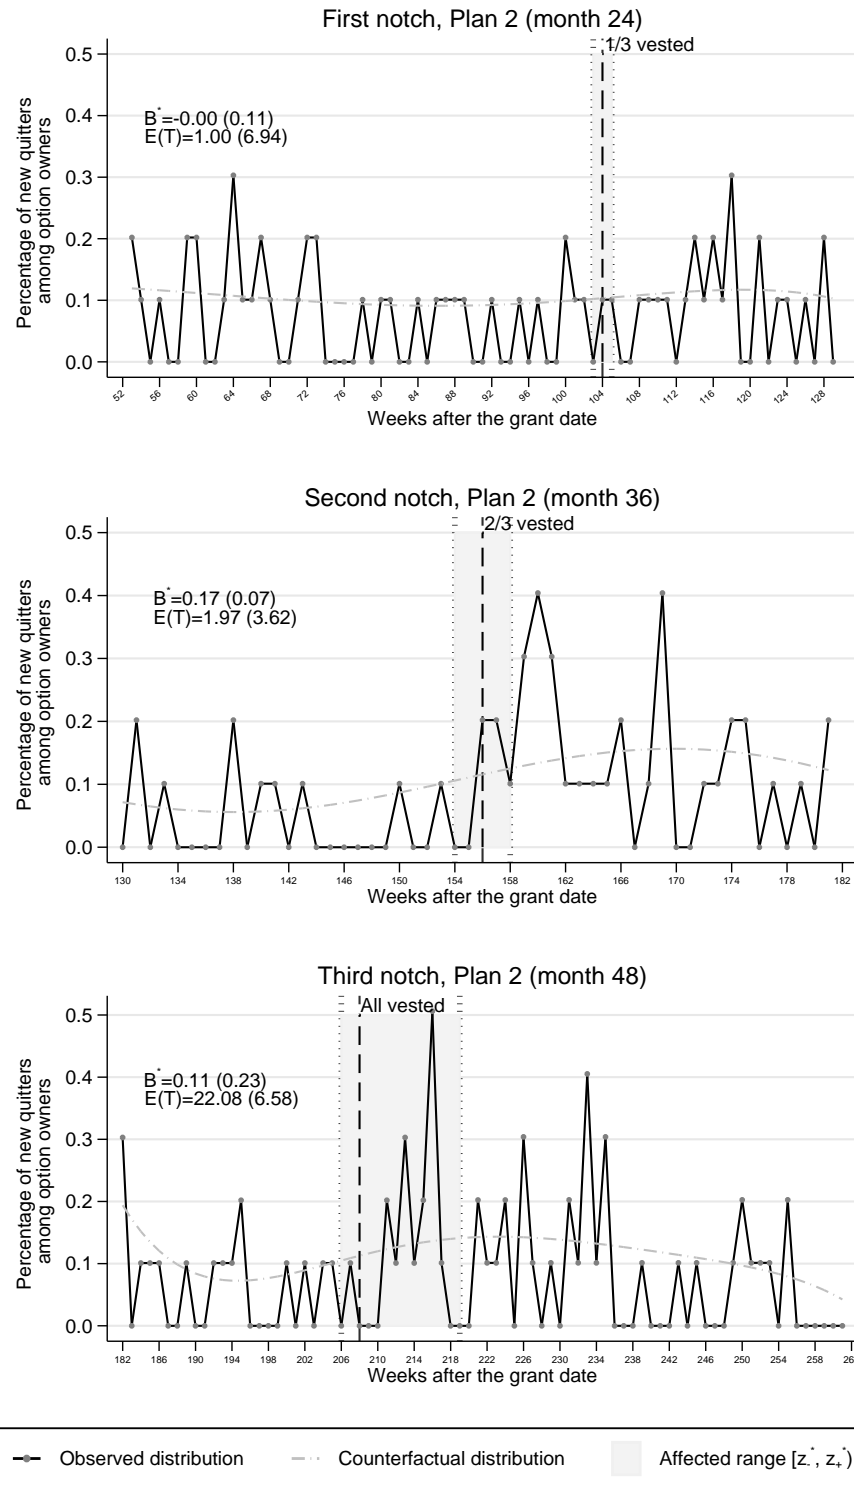


Figure A3: Estimation of bunching using weekly turnover rates (Plan 2)

NOTES: The horizontal axes show months after option granting, with the grant date being the first day of month 0. The three panels show bunching at the three vesting dates for Plan 2. The solid lines trace out the observed distribution of quitters. The dashed lines show the counterfactual distributions. The shaded area indicates the affected range. Sizes of bunching at the three vesting dates, measured by excess mass  $B^*$ , are reported at the top-left corner of each plot. The corresponding standard errors are in parentheses. Also reported are the expected number of months an option owner is retained,  $E(T)$ .