Online Appendix to

DEADLINES VERSUS CONTINUOUS INCENTIVES: EVIDENCE FROM THE PATENT OFFICE

By: Michael D. Frakes and Melissa F. Wasserman

Annual Time-Trend in End-Loading Rates: Regression Analysis

Figure 1 of the text plots a time trend in the average rates of first-round review end-loading (incidence of first-round review being completed on the final day of the quota period) by year. In this section of the Online Appendix, we demonstrate that this trend cannot be explained by annual fluctuations in certain factors. For these purposes, we estimate the following regression specification:

$$End - Loaded_{at} = \alpha + \boldsymbol{\gamma}_j + \sum_{t=2003}^{t=2015} \boldsymbol{\beta}_{\mathbf{t}} \boldsymbol{Year}_{\mathbf{t}} + \boldsymbol{\Theta}\mathbf{X}_{at} + \varepsilon_{ajt} \tag{A1}$$

Where a indicates the first-round review associated with the individual application, j indicates the Art Unit of the associated examiner and t indicates the year in which the focal first-round review was completed. $\mathbf{Year_{ait}}$ is a set of dummy variables indicating the first-round review year. We include a set of Art Unit fixed effects, $\boldsymbol{\gamma}_j$, to address possible concerns that Figure 1 can be explained by fixed differences in end-loading tendencies across Art Units combined with large swings around 2011 in the composition of Art Units receiving applications. We also include a set of application controls in $\mathbf{X}_{at}$ to address similar concerns that swings in the types of applications filed around 2011 (combined with inherent tendencies for some types of applications to be end-loaded) could explain the drop in end-loading rates around 2011. Application controls include an indicator for the incidence of a small-entity applicant, the number of independent claims, the number of words in the independent claims, the average number of words in independent claims and the minimum number of words in a claim across the independent claims.

As demonstrated by the following figures, we find that the pattern presented in Table 1 is completely robust to these controls. We do this two ways: (1) showing first-round reviews completed by GS-14 examiners only and (2) pooling across all examiners (and including examiner GS-level fixed effects).

Figure A1: End-Loading Rates by Year, Controlling for Application Characteristics, Examiner Art Unit Effects and Examiner GS-Level Effects (All Examiners Pooled)

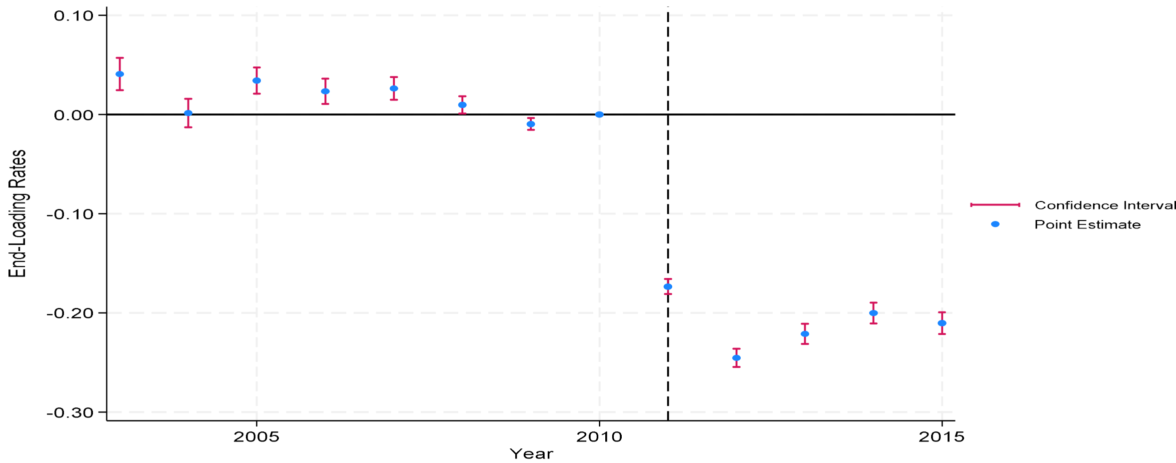

Notes: this figure plots the estimated coefficients of the year effects from specification (A1), effectively plotting an annual time trend in end-loading rates across first-round reviews (incidence of first-round reviews being completed on the last quota day), while controlling for various application characteristics, Art Unit effects and examiner GS-level effects. N = 3,573,462.

Figure A2: End-Loading Rates by Year, Controlling for Application Characteristics, and Examiner Art Unit Effects (GS-14 Examiners)

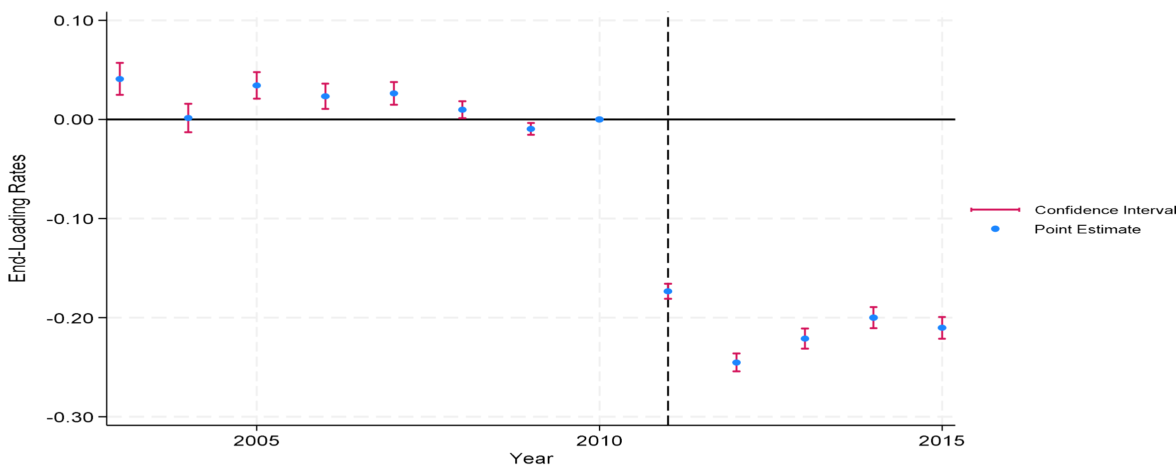

Notes: this figure plots the estimated coefficients of the year effects from specification (A1), effectively plotting an annual time trend in end-loading rates across first-round reviews (incidence of first-round reviews being completed on the last quota day), while controlling for various application characteristics and Art Unit effects. Only GS-14 examiners are included. N = 1,590,175.

Smoothing Analysis over Day-by-Day of Bi-Week: Randomization Inference

Figure 5 of the text shows mean rates of various outcomes associated with first-round reviews across the given days of the bi-week period, including the rate by which the first-round review is allowed. In the pre-reform period, we document a strongly negative trend over the bi-week in first-round allowance rates. To demonstrate how much of an outlier the actual first-round review allowance patterns are across the bi-week, we simulate 100 different placebo first-round review patterns. In each simulation, we randomly assign each first-round review a placebo bi-week day and then calculate the average first-round-review allowance rate across the relevant placebo bi-week days. Doing so across each simulation, one can see in Figure A3 that the actual bi-week pattern of first-round allowance rates—and particularly its slope—is an extreme outlier relative to the set of placebo bi-week patterns. In Figure A3, we depict the actual pattern in the black line and the placebo patterns in the faint gray lines. This figure suggests that chance alone simply cannot explain the striking first-round allowance pattern produced in Panel A of Figure 5.

In Figure A4, we conduct the same exercise but now focusing on the post-reform period. As demonstrated, now the actual first-round-review allowance-rate pattern over the bi-week is much more aligned with that of the set of placebo patterns. This is consistent with the near total elimination of deadline effects in the post-reform period.

Figure A3: Relationship between Bi-Week Days and First-Round Allowance Rates in Pre-Reform Period, Randomization Inference (Actual Pattern versus 100 Simulated Placebo Patterns)

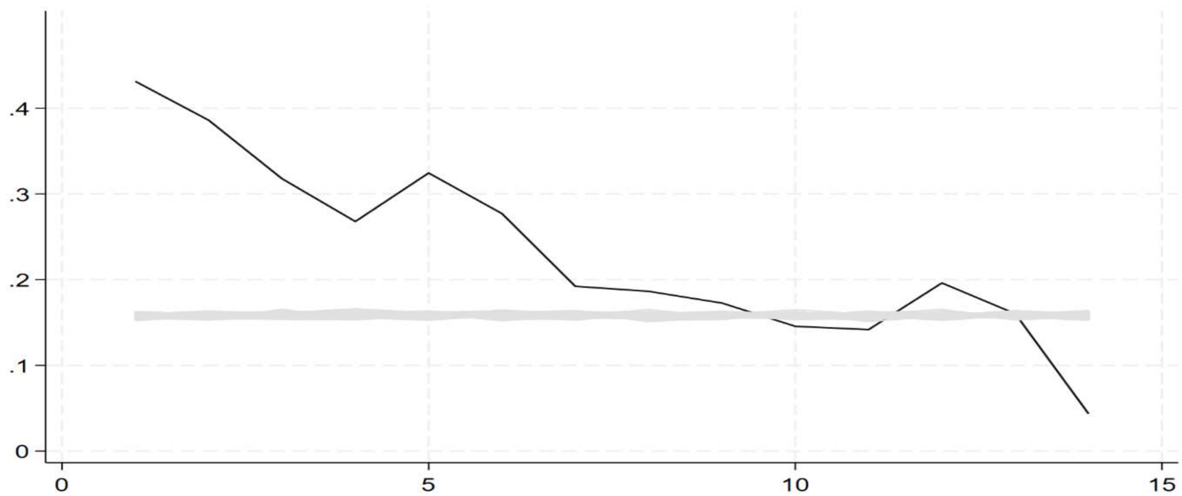

Notes: the black line reproduces the pre-reform results from Panel A of Figure 5, depicting the mean first-round allowance rates by day of the bi-week in the 2008-2010 period for GS-14 examiners. N = 341423. The gray lines represent 100 placebo trends in first-round allowance rates over the bi-week (after randomly generating placebo bi-week days).

Figure A4: Relationship between Bi-Week Days and First-Round Allowance Rates in Post-Reform Period, Randomization Inference (Actual Pattern versus 100 Simulated Placebo Patterns)

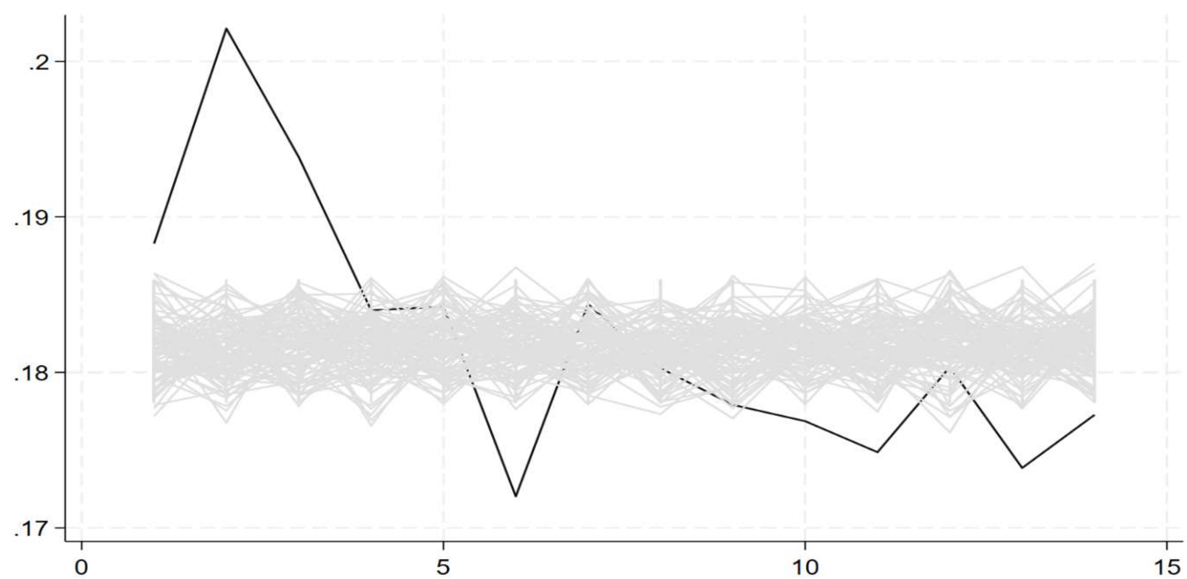

Notes: the black line reproduces the pre-reform results from Panel A of Figure 5, depicting the mean first-round allowance rates by day of the bi-week in the 2012-2014 period for GS-14 examiners. N = 536913. The gray lines represent 100 placebo trends in first-round allowance rates over the bi-week (after randomly generating placebo bi-week days).

Smoothing Analysis over Day-by-Day of Bi-Week: Weighted Figures

In Figure A5, we present a version of Figure 5 from the text that weights each bi-week point by the number of reviews associated with that point. What this effectively allows us to do is to combine the results from Figures 3 and 5, which, in turn, allows us to demonstrate smoothing of two categories of decisions: (1) smoothing of the timing of the review completions themselves (less of a concentration of completions on the final day) and (2) smoothing of the outcomes of those reviews (e.g., smoothing of first-round allowance rates over the bi-week).

FIGURE A5: ESTIMATED RELATIONSHIP BETWEEN THE INDICATED OUTCOMES ASSOCIATED WITH THE FIRST-ROUND DECISIONS AND THE INCIDENCE OF THAT FIRST-ROUND DECISION OCCURRING ACROSS THE GENERALIZED BI-WEEKLY PERIOD, BEFORE AND AFTER 2011 PATENT OFFICE PERFORMANCE APPRAISAL REFORM (GS-14 EXAMINERS), WITH APPLICATION-COUNT WEIGHTS

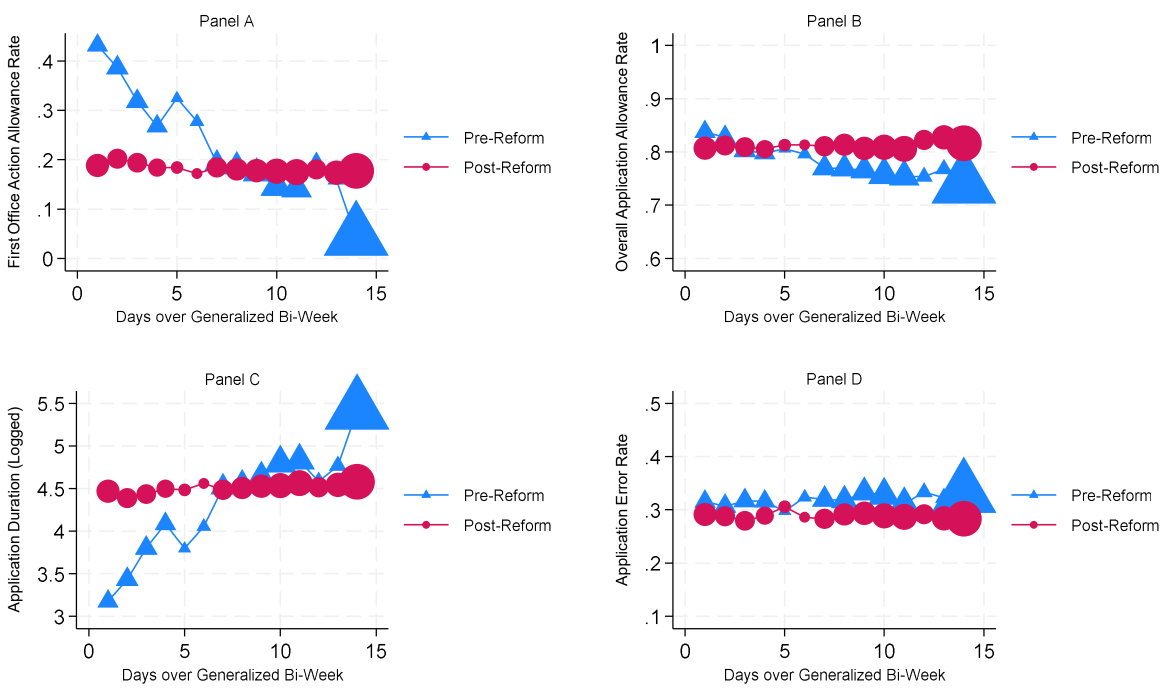

Notes: this figure replicates that of Figure 5 except that it weights each point (representing each day of the bi-week) by the number of reviews associated with that day of the bi-week.

Smoothing Analysis over Day-by-Day of Bi-Week: Leads Analysis

Figure 5 of the text groups the pre-reform period into a three-year group (2008-2010) and the post-reform period into a three-year group (2012-2014). It demonstrates a strong smoothing in connection with the reform in various decisions and outcomes associated with first-round reviews. In this section of the Online Appendix, we demonstrate that the smoothing commences at the moment of the reform and not prior to the reform. For the sake of brevity, we show this for the two outcomes in which we see the largest smoothing response in Figure 5: first-round allowance and application duration (recall that overall allowance rates and EPO allowance rates are already smoothed in the pre-period, consistent with the shotgun rejection theory). Only in 2011 do we see the smoothing start—recall that the reform occurred in the third quarter of the 2011 fiscal year (between April and June of 2011). The patterns are remarkably consistent year-by-year in the pre- and post-reform periods.

FIGURE A6: ESTIMATED RELATIONSHIP BETWEEN THE INDICATED OUTCOMES ASSOCIATED WITH THE FIRST-ROUND DECISIONS AND THE INCIDENCE OF THAT FIRST-ROUND DECISION OCCURRING ACROSS THE GENERALIZED BI-WEEKLY PERIOD, YEAR-BY-YEAR LEADING UP TO AND FOLLOWING THE 2011 PATENT OFFICE PERFORMANCE APPRAISAL REFORM (GS-14 EXAMINERS)

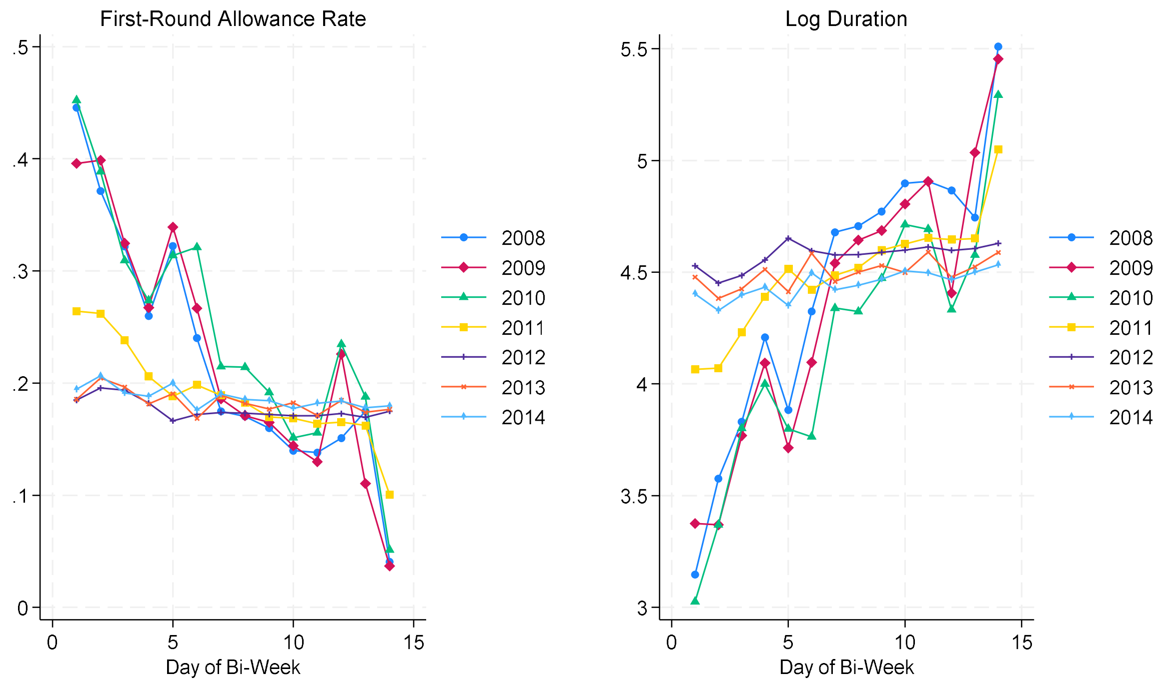

Smoothing Analysis over Day-by-Day of Bi-Week: Regression Analysis

In Figure 5 of the text, we break down first-round review decisions into the day of the bi-week in which the reviews were completed, where 1 indicates the first day of the quota period and 14 indicates the final day of the quota period, and then show mean rates of the following outcomes associated with those first-round reviews: (1) allowance on the first-round review, (2) ultimate allowance of the application (3) log of duration in days of the ultimate application and (4) the incidence of an examination error (based on EPO benchmarking). Moreover, we show these mean outcomes separately for the 2008-2010 (calendar-year) pre-reform period and the 2012-2014 (calendar-year) post-reform period (excluding 2011 as the reform occurred in the middle of the calendar year).

In this section of the Online Appendix, we now turn to a regression implementation of that analysis, whereby for each of the pre- and post-reform period we estimate the following specification across the various outcomes explored in Figure 5, where the unit of observation is a given first-round review:

$$Outcome_{ait} = \alpha + \boldsymbol{\gamma}_{\mathbf{i}} + \sum_{t=1}^{t=14} \boldsymbol{\beta}_{\mathbf{t}}\mathbf{Quota-Day_{ait}} + \boldsymbol{\theta}\mathbf{X}_{ai} + \varepsilon_{ait} \tag{A2}$$

Where a indicates the first-round review associated with the individual application, i indicates the associated examiner and t indicates the day of the bi-week quota period in which the focal first-round review was completed. $\mathbf{Quota-Day_{ait}}$ is a set of dummy variables indicating the day of the bi-week in which the first-round review was completed.

A set of examiner fixed effects is included in $\boldsymbol{\gamma}_{\mathbf{i}}$. This allows us to account for potential concerns associated with Figure 5 in which those examiners that tend to end-load their reviews may have fundamentally different levels of the relevant outcome variables. By including examiner fixed effects, we look within given examiners and ask how first-round review outcomes are a function of the day in which the first-round review is completed.

A set of application-level characteristics is included in $\mathbf{X}_{ai}$. These include an indicator for whether the applicant is a small-entity (as the Patent Office classifies such term for the purposes of fee-setting) and controls for the following claim characteristics as of the time of the publication of the underlying patent application: number of independent claims, word count of independent claims, average number of words in independent claims, and minimum number of words across the independent claims.

As discussed in the text, an advantage of presenting only means in Figure 5 is that it allows for the possibility that the smoothing in application outcomes we observe over the bi-week period in connection with the reform could be a function of both, or either of, (i) a reduction in time-mismanagement behaviors (e.g., procrastination) and (ii) a reduction in application sorting over the bi-week (e.g., reserving close-call reviews for the end of the bi-week and more straightforward reviews for the beginning of the bi-week). After all, if we had controlled for characteristics that would form the basis of the task-sorting behavior, we would abstract away from such behavior. What is nonetheless helpful about including application controls is that it allows us to observe whether the patterns documented in Figure 5 persist even after we control for factors that could form the basis of a task-sorting story—i.e., it allays concerns that Figure 5 is predominantly picking up a task-sorting story.

As demonstrated by the following figures—Figures A7 – A10, the pre-reform patterns look essentially the same whether we include these controls or whether we simply show means. Moreover, the flattening out following the reform also appears unaffected by the inclusion of such controls. Altogether, it does not appear that a sorting story based on the included application characteristics—including claim-based characteristics that speak to the scope of the requested claims—can account for patterns presented in Figure 5 in the text. This is suggestive evidence that time mismanagement behaviors—e.g., procrastination—may be accounting for the pre-reform deadline patterns and that the 2011 reform may be unwinding such behaviors.

Figure A7: Relationship between First-Round Review Allowance Rates and Day of Bi-Week, Pre- and Post-Reform, Controlling for Application Characteristics and Examiner Fixed Effects

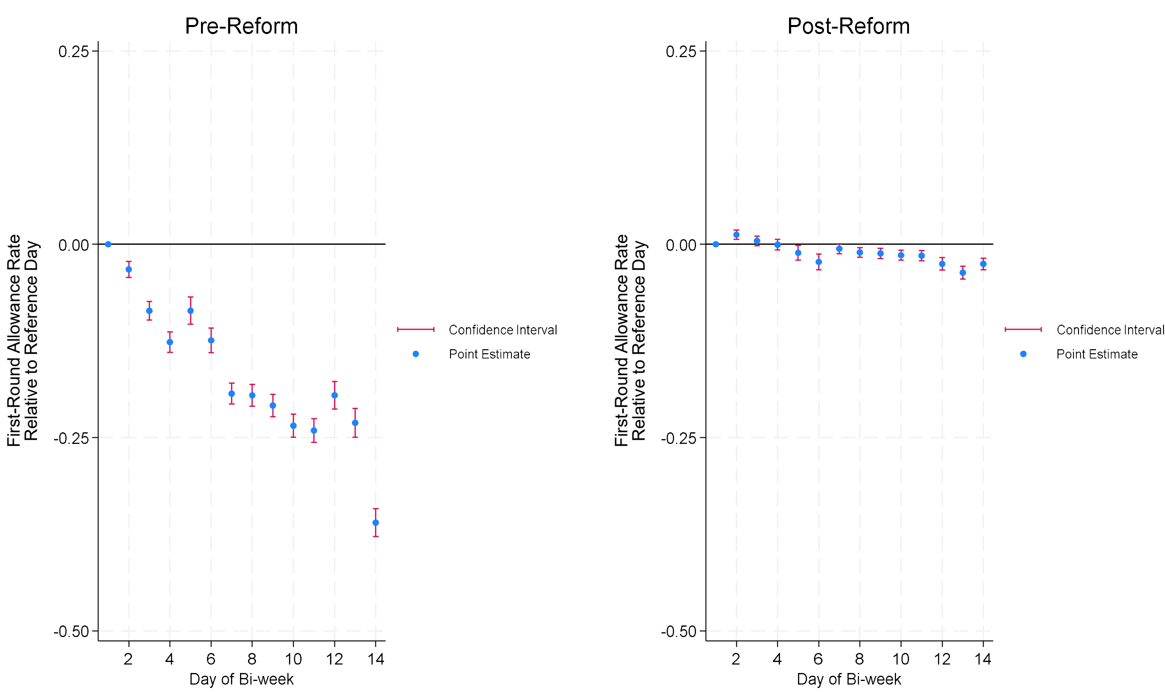

Notes: the number of observations in the pre-reform specification is 322091 and the number of observations in the post-reform specification is 497473. Each figure presents the estimated $\boldsymbol{\beta}_t$ coefficients from specification A2.

Figure A8: Relationship between Overall Allowance Rate and Day of Bi-Week on Which First-Round Review was Completed, Pre- and Post-Reform, Controlling for Application Characteristics and Examiner Fixed Effects

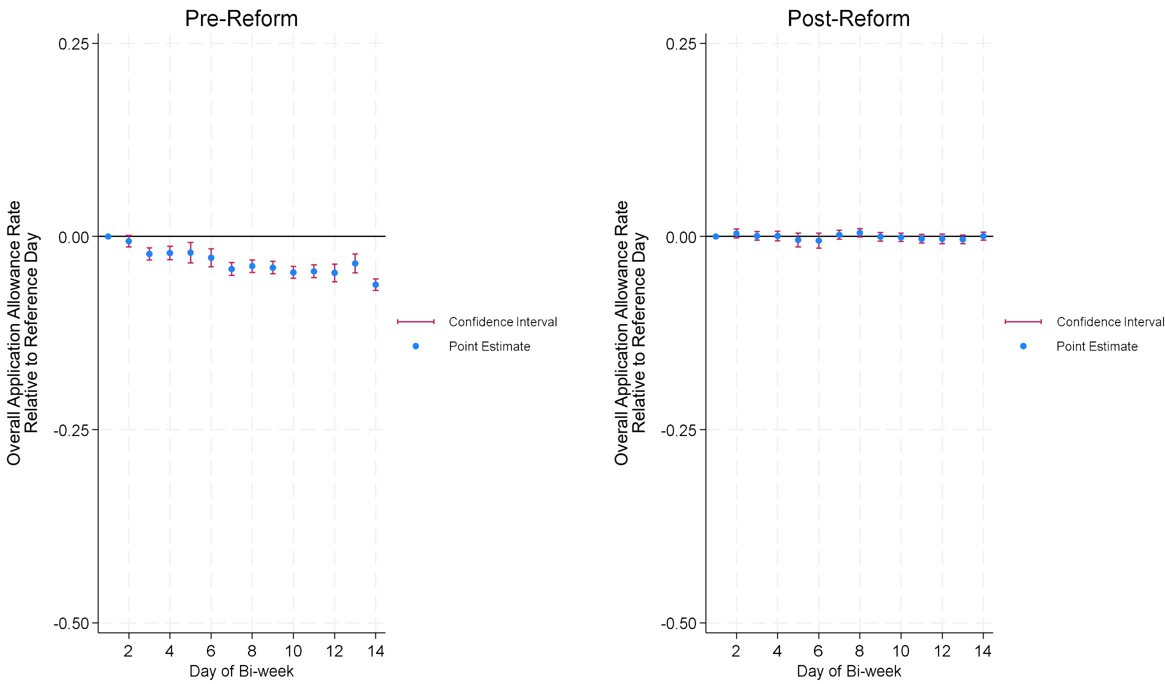

Notes: the number of observations in the pre-reform specification is 322091 and the number of observations in the post-reform specification is 497473. Each figure presents the estimated $\boldsymbol{\beta_t}$ coefficients from specification A2.

Figure A9: Relationship between Log of Application Duration and Day of Bi-Week on Which First-Round Review was Completed, Pre- and Post-Reform, Controlling for Application Characteristics and Examiner Fixed Effects

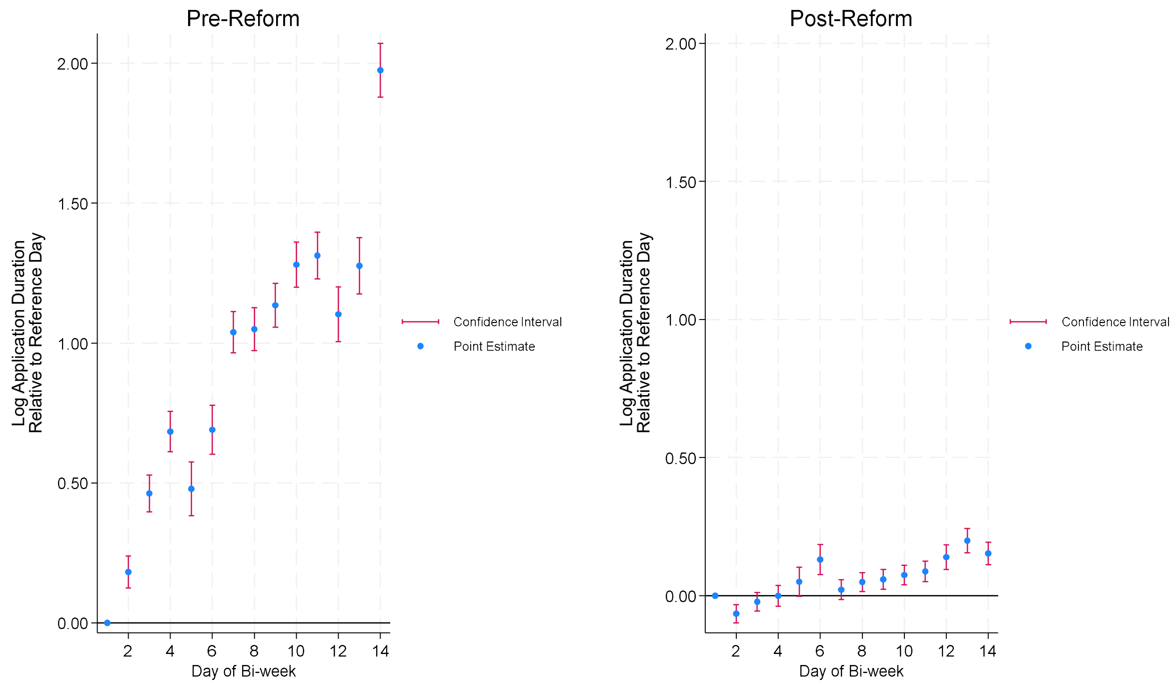

Notes: the number of observations in the pre-reform specification is 322091 and the number of observations in the post-reform specification is 497473. Each figure presents the estimated $\boldsymbol{\beta_t}$ coefficients from specification A2.

FIGURE A10: RELATIONSHIP BETWEEN INCIDENCE OF APPLICATION ERROR AND DAY OF BI-WEEK ON WHICH FIRST-ROUND REVIEW OF U.S. APPLICATION WAS COMPLETED, PRE- AND POST-REFORM, CONTROLLING FOR APPLICATION CHARACTERISTICS AND EXAMINER FIXED EFFECTS

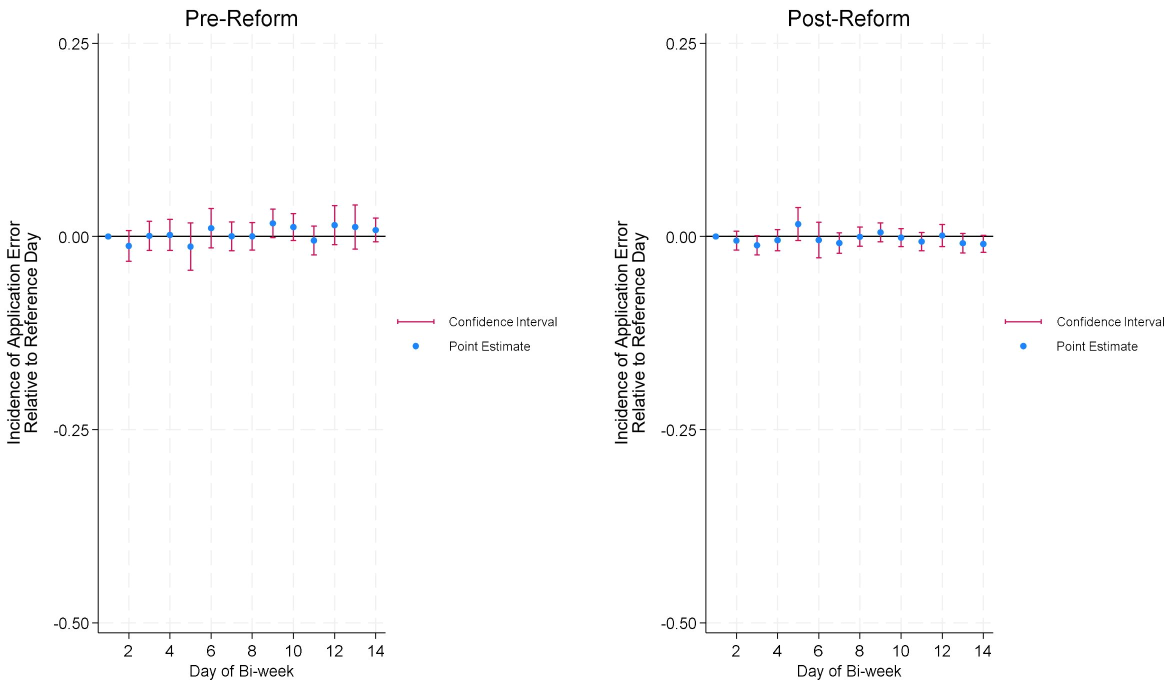

Notes: the number of observations in the pre-reform specification is 91737 and the number of observations in the post-reform specification is 126322 (smaller than in Figures A7 – A9 given that this figure is based on a subset of applications associated with U.S.-issued patents whose underlying innovations are also the subjects of applications at the EPO). Each figure presents the estimated $\boldsymbol{\beta}_t$ coefficients from specification A2.

Smoothing Analysis over Day-by-Day of Bi-Week: Sorting

We now assess whether examiners appear to (1) sort applications across the bi-week based on observable application characteristics and (2) reduce such sorting following the 2011 Patent Office Reform that imposes continuous temporal incentives. In particular, we explore mean rates of the following claim characteristics across the days of the bi-week, pre- and post-reform: (1) number of independent claims, (2) total number of words in independent claims, (3) minimum number of words across independent claims and (4) predicted likelihood of allowance on first-round review, based on a regression of first-round allowance on a full set of application characteristics.[1]

As demonstrated by Figure A11, we do document a degree of sorting on these observable characteristics across the bi-week, with examiners starting with applications with arguably narrower claim scope, as proxied by the total number of words across independent claims and the minimum number of words per claim across the independent claims. As demonstrated by Panel A, examiners begin the bi-week with applications with characteristics that predict a higher degree of allowance on the first-office action.

Moreover, Figure A11 also demonstrates that this sorting based on observable application characteristics appears to go away following the 2011 reform.

Finally, Figure A12 situates the predicted first-round allowance rates patterns (Panel A) against the actual first-round allowance rate patterns (Panel B, replicating the results from Figure 5). We place the y-axes on the same scale across these two panels. Despite the task sorting implied by Figure A11, Figure A12 demonstrates that the degree to which first-round application decisions vary across the bi-week is substantially more pronounced than the degree to which predicted first-round allowances vary over the bi-week. This suggests that whatever task sorting is present, it likely pales in comparison with the other behavioral mechanisms leading to the pre-reform deadline effects documented in Figures 1-5.

[1] Characteristics include an indicator for the applicant being a small entity and on a series of claim characteristics (and their squares): number of independent claims, number of dependent claims, total number of words in independent claims, total number of words in dependent claims, average number of words in independent claims, average number of words in dependent claims, minimum word count among word counts in independent claims, and minimum word count among word count in dependent claims.

FIGURE A11: RELATIONSHIP BETWEEN THE INDICATED OUTCOMES ASSOCIATED WITH THE FIRST-ROUND DECISIONS AND THE INCIDENCE OF THAT FIRST-ROUND DECISION OCCURRING ACROSS THE GENERALIZED BI-WEEKLY PERIOD, BEFORE AND AFTER 2011 PATENT OFFICE PERFORMANCE APPRAISAL REFORM (GS-14 EXAMINERS)

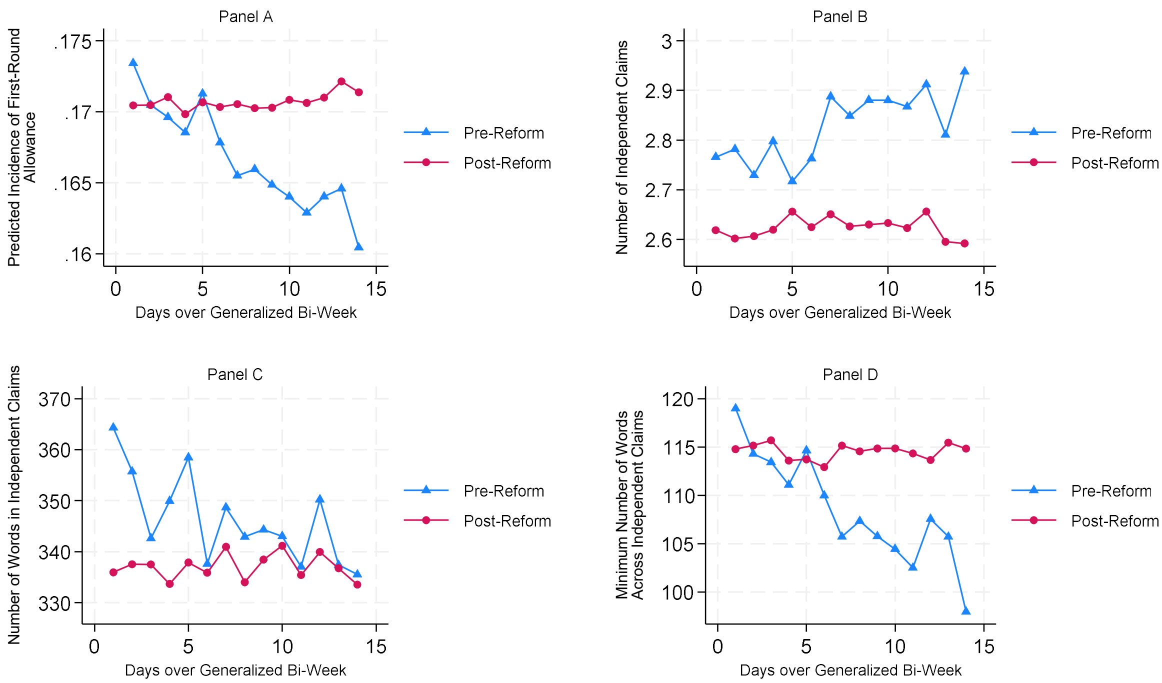

Notes: this Figure repliactes that of Figure 5 except focusing on outcomes reflective of application characteristics.

FIGURE A12: PREDICTED FOAM ALLOWANCE RATES BY DAY OF BI-WEEK (PRE- AND POST-REFORM) VERSUS ACTUAL FOAM ALLOWANCE RATES BY DAY OF BI-WEEK (GS-14 EXAMINERS)

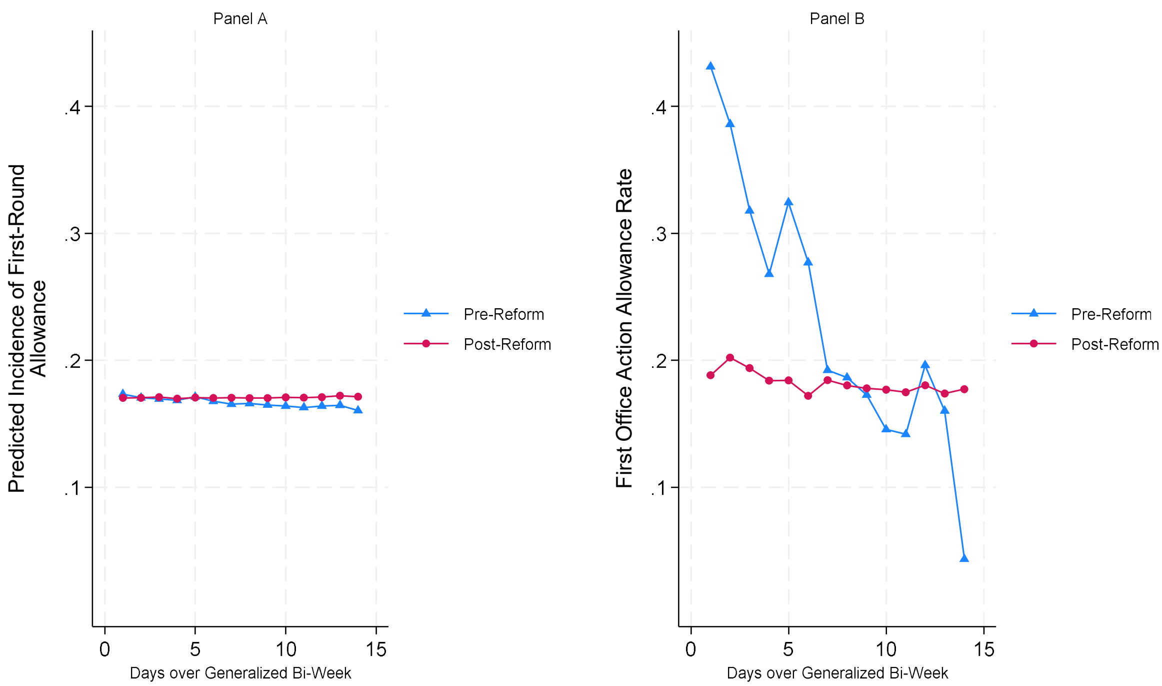

Notes: each Panel of this Figure repliactes that of Figure 5 except focusing on the indicated outcomes. Panel A presents results for predicted FOAM allowance rates (predicted based on application characteristics), whereas Panel B focuses on actual FOAM allowance rates.

Duration Results in Levels

FIGURE A13: ESTIMATED RELATIONSHIP BETWEEN APPLICATION DURATION IN DAYS AND THE INCIDENCE OF THE APPLICATION'S INITIAL REVIEW OCCURRING ON THE LAST DAY OF THE QUOTA PERIOD, BY YEAR (GS-14 EXAMINERS)

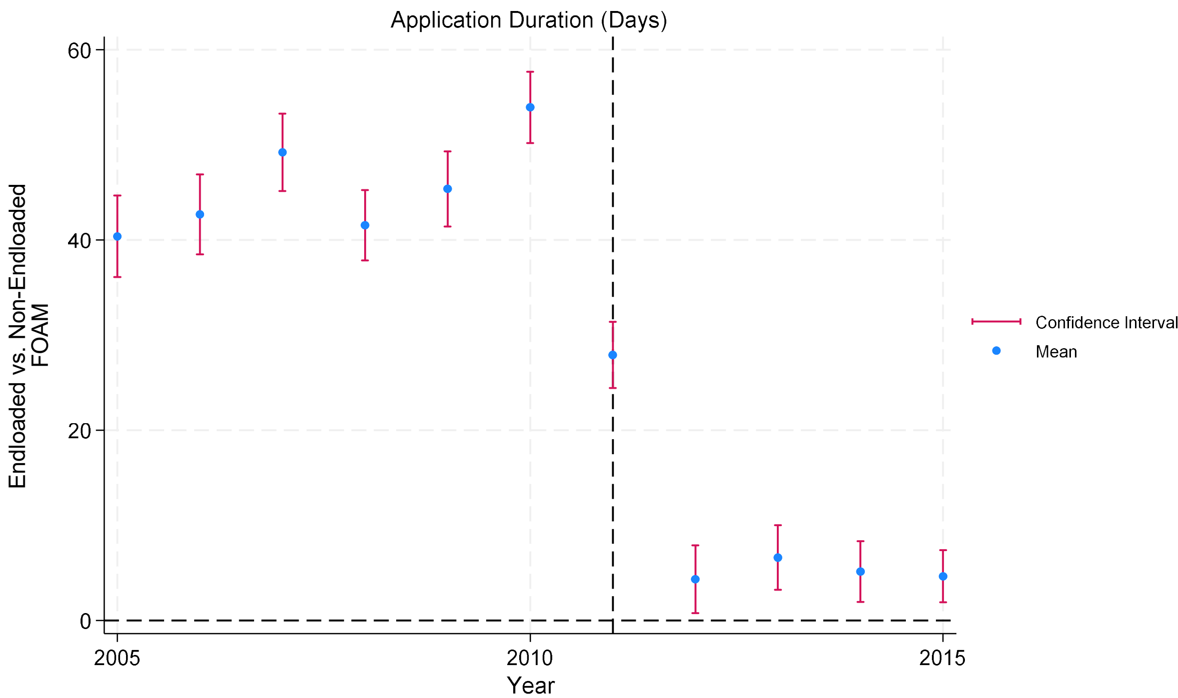

Notes: this figure replicates the duration figure from Figure 4 except using duration in levels and not logs.

FIGURE A14: RELATIONSHIP BETWEEN APPLICATION DURATION IN DAYS AND THE INCIDENCE OF THE APPLICATION'S FIRST-ROUND DECISION OCCURRING ACROSS THE GENERALIZED BI-WEEKLY PERIOD, BEFORE AND AFTER 2011 PATENT OFFICE PERFORMANCE APPRAISAL REFORM (GS-14 EXAMINERS)

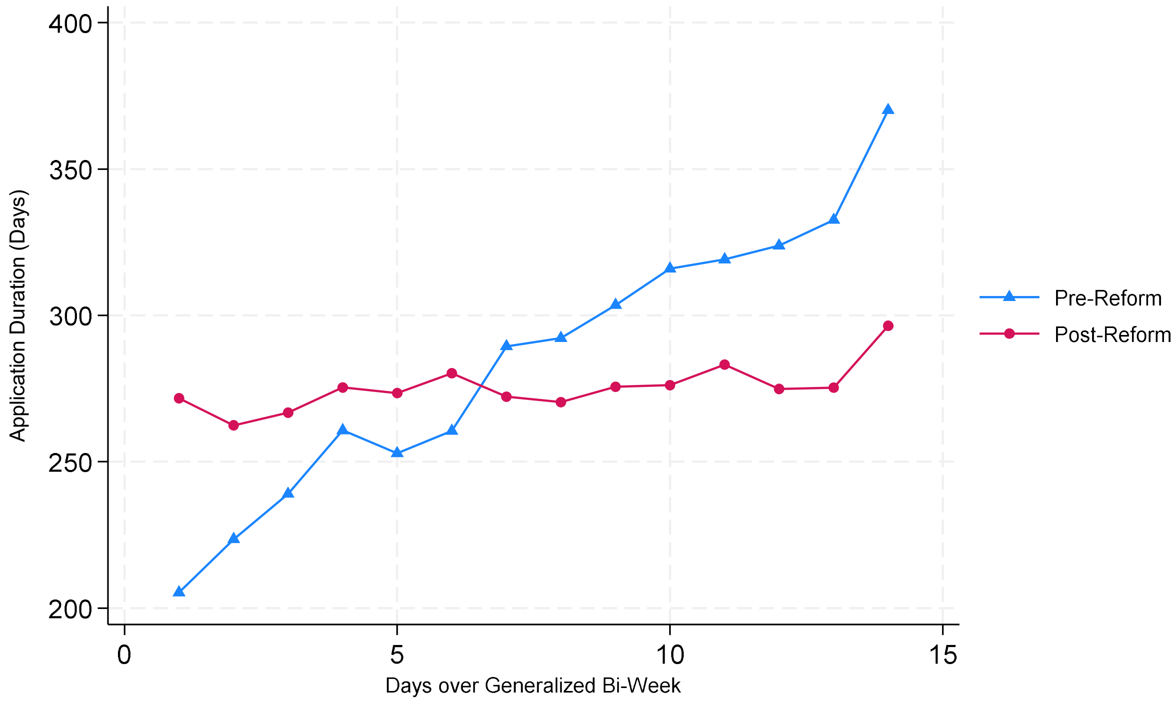

Notes: this figure replicates the duration figure from Figure 5 except using duration in levels and not logs.

Second-Office-Action Non-Final-Rejection Analysis

Our primary markers of shotgun rejections—summarily rejecting in the face of deadline-looming time crunches—include high rejection rates on first-round reviews at the end of the bi-week, long overall application durations for first-round reviews completed at the end-of the bi-week, and overall similar allowance rates for applications whose first-round reviews were and were not completed at the end of the bi-week. In this section of the Online Appendix, we consider the incidence of yet another marker of a shotgun rejection: second office-action non-final rejections.

Recall from Part II of the text that rejections on second-round reviews are typically "final." When examiners issue "non-final" rejections in the second round, they are generally introducing new bases of rejections and are thus implicitly admitting that they performed an inadequate review on the first round. In this section, we replicate the analyses underlying Figures 4, 5, and 7 of the text but now using the incidence of a second office-action non-final rejection as the outcome variable.

We begin in Figure A15 with the counterpart to Figure 3.

FIGURE A15. RELATIONSHIP BETWEEN FOAM END-LOADING AND INCIDENCE OF SECOND OFFICE ACTION NON-FINAL REJECTION, BY YEARS BEFORE AND AFTER 2011 PATENT OFFICE PERFORMANCE APPRAISAL REFORM (GS-14 EXAMINERS)

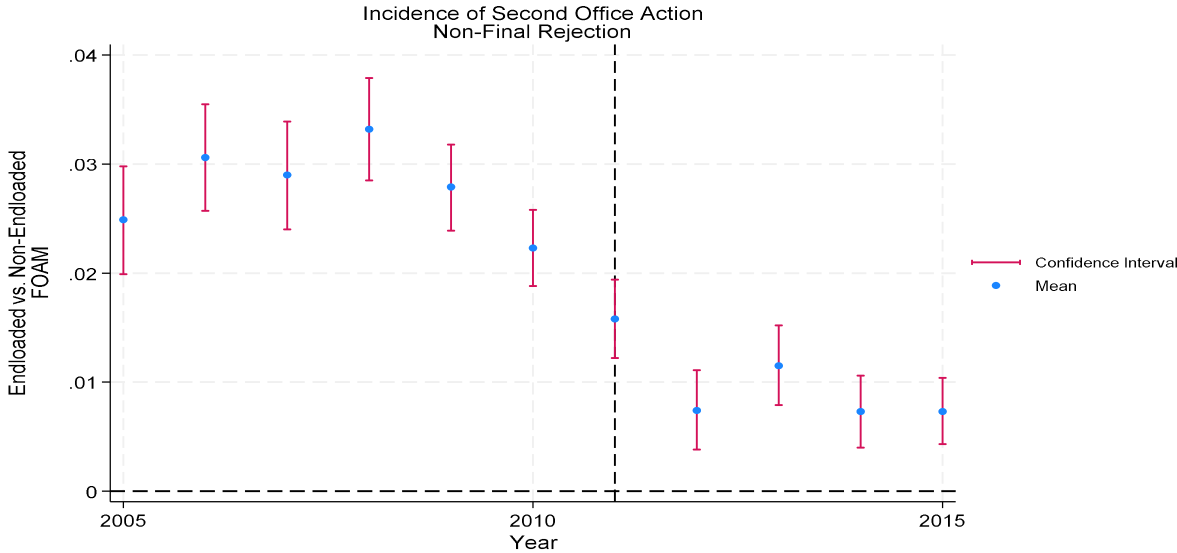

Notes: N ranges from 107,887 in 2005 to 226,293 in 2015. Each estimate is from a separate annual regression, where the unit of observation is a given first-round review and where the incidence of a second office action non-final rejection associated with the focal review is regressed on an indicator for whether the first-round review was completed on the final day of the quota period. The estimated coefficient of this indicator is presented. The regressions also control for examiner fixed effects and various time-varying characteristics of the examiner (GS-level and experience fixed effects) in addition to an indicator for the application being associated with a small entity applicant.

FIGURE A16: RELATIONSHIP BETWEEN THE INCIDENCE OF A SECOND OFFICE ACTION NON-FINAL REJECTION AND THE INCIDENCE OF THAT FIRST-ROUND DECISION OCCURRING ACROSS THE GENERALIZED BI-WEEKLY PERIOD, BEFORE AND AFTER 2011 PATENT OFFICE PERFORMANCE APPRAISAL REFORM (GS-14 EXAMINERS)

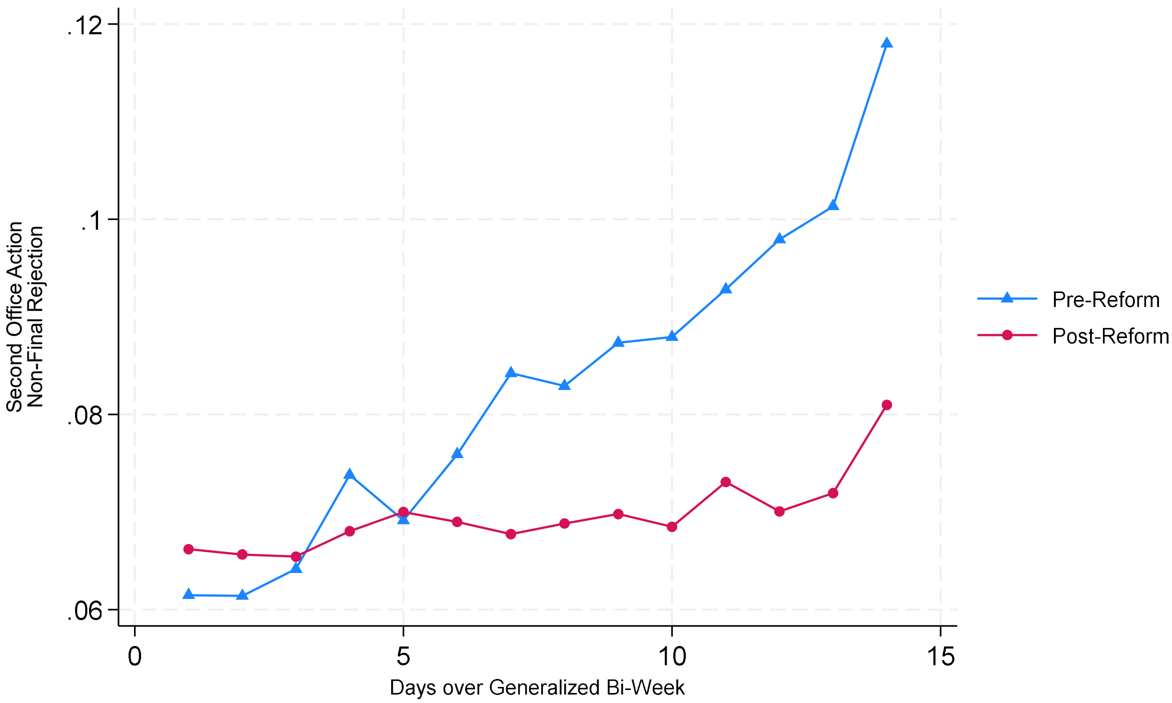

Notes: In this figure, we plot the mean incidence of second-office-action non-final rejections by the date of the bi-week in which the associated first-round review was completed, separately for the pre-reform period (2008-2010, N = 375,340) and the post-reform period (2012-2014, N=546,764). Only applications reviewed by GS-14 examiners are included.

FIGURE A17: DYNAMIC SECOND-OFFICE ACTION NON-FINAL REJECTION RESULTS FOR GS-14 RELATIVE GS-11-AND-BELOW EXAMINERS

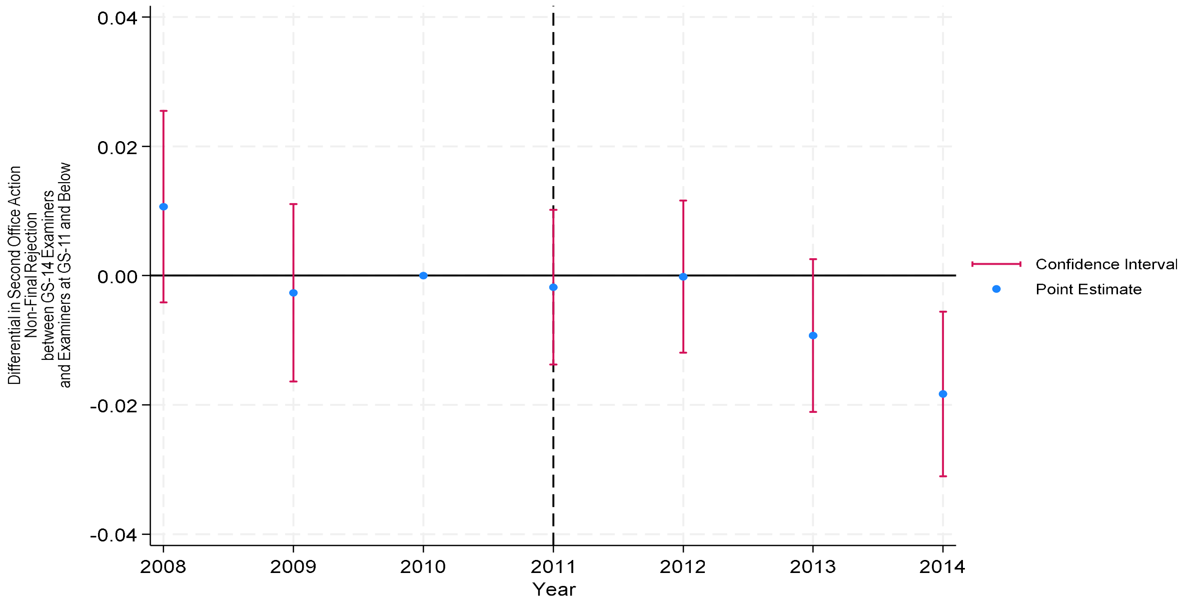

Notes: this figure replicates that of Figure 7 in the text except using as the outcome variable the incidence of a second office-action non-final rejection.

Application Review Error Analysis

In the text, we explore the relationship between the days of the bi-week in which examiners review the first office actions and the error rate associated with the overall applications. An error is coded as having occurred if the application is associated with an allowance at the U.S. Patent Office and a rejection at the EPO (Type I error) or if the application is associated with a rejection at the U.S. Patent Office and an acceptance at the EPO (Type II error). We find no relationship between the timing of the first-round review (including the end-loading of the first-round review) and the ultimate application error rate. Moreover, we find no overall effect of the 2011 reform on the application error rate (via the difference-in-difference analysis we conduct in Part IV). This is consistent with a story in which, pre-reform, examiners issue shotgun rejections leading up to the end of the quota period. After all, the purpose of the shotgun rejection behavior is to buy time in later rounds to ultimately give the application an accurate review, and these accuracy measures are based on the final examination outcome.

In this part of the Online Appendix, we perform the analysis underlying Figures 4, 5 and 7 of the text separately for Type I and Type II errors. As we demonstrate, the overall error-rate results from the text replicate across the separate error types.

FIGURE A18: ESTIMATED RELATIONSHIP BETWEEN INCIDENCE OF TYPE I AND TYPE II APPLICATION ERRORS AND THE INCIDENCE OF THE FIRST-ROUND REVIEW OCCURRING ON THE LAST DAY OF THE QUOTA PERIOD, BY YEAR (GS-14 EXAMINERS)

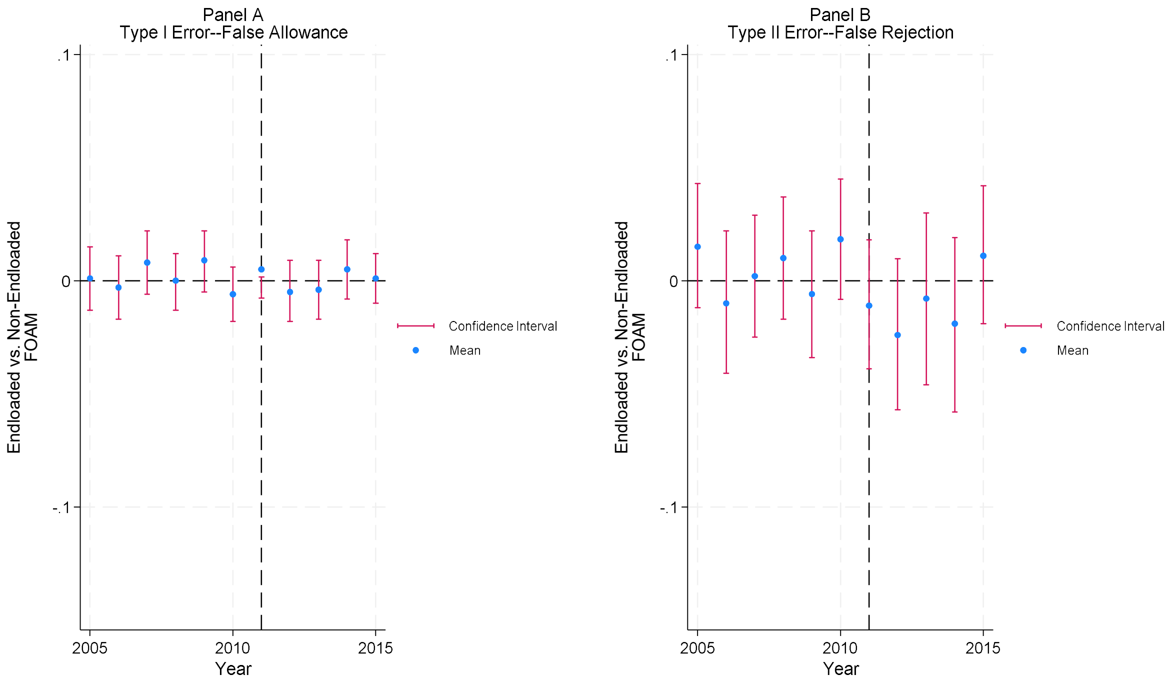

Notes: this figure replicates Figure 4 Panel D but separately for Type I and Type II examination errors. N = 24749 (2005) – 40484 (2015) for Panel A. N = 6619 (2005) to 9406 (2015) for Panel B.

FIGURE A19: RELATIONSHIP BETWEEN TYPE I AND TYPE II ERRORS AND THE INCIDENCE OF THE ASSOCIATED FIRST-ROUND DECISION OCCURRING ACROSS THE GENERALIZED BI-WEEKLY PERIOD, BEFORE AND AFTER 2011 PATENT OFFICE PERFORMANCE APPRAISAL REFORM (GS-14 EXAMINERS)

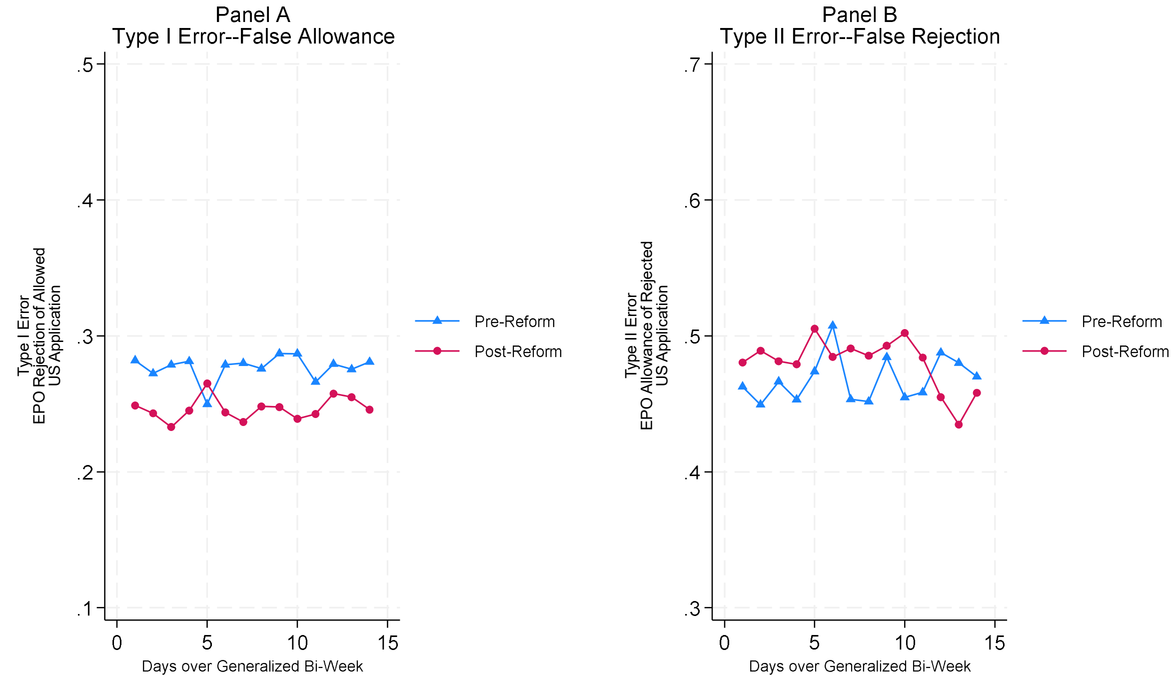

Notes: this figure replicates Panel D of Figure 5 but separately for Type I and Type II errors. For Panel B, N = 21145 in the pre-period and 23802 in the post period. For Panel A, N = 77750 in the pre period and 107180 in the post period.

FIGURE A20: DYNAMIC APPLICATION REVIEW ERROR RATE RESULTS FOR GS-14 RELATIVE GS-11-AND-BELOW EXAMINERS, SEPARATELY BY TYPE I AND TYPE II ERRORS

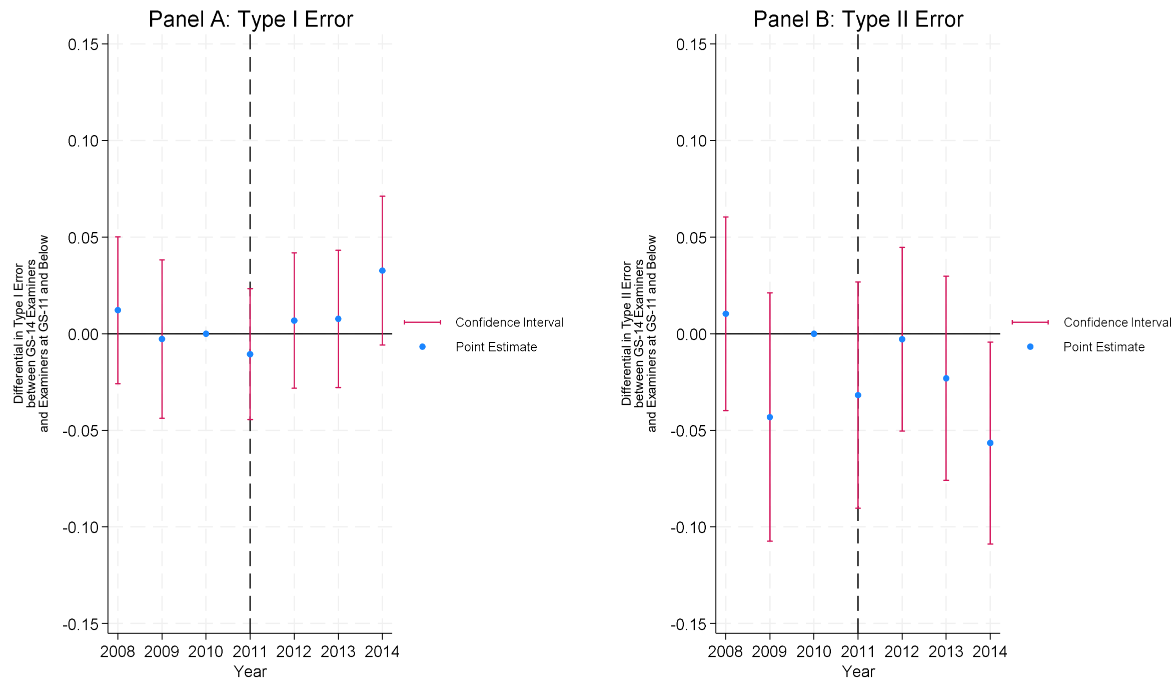

Notes: this figure replicates Panel C of Figure 7 but separately for Type I and Type II errors. For Panel B, N = 109932 and for Panel A, N = 297212.

Overall Reform-Effects Analysis: Additional Claim-Characteristic Controls

FIGURE A21: DYNAMIC EXAMINATION QUALITY RESULTS FOR GS-14 RELATIVE GS-11-AND-BELOW EXAMINERS, WITH CLAIM CHARACTERISTIC CONTROLS

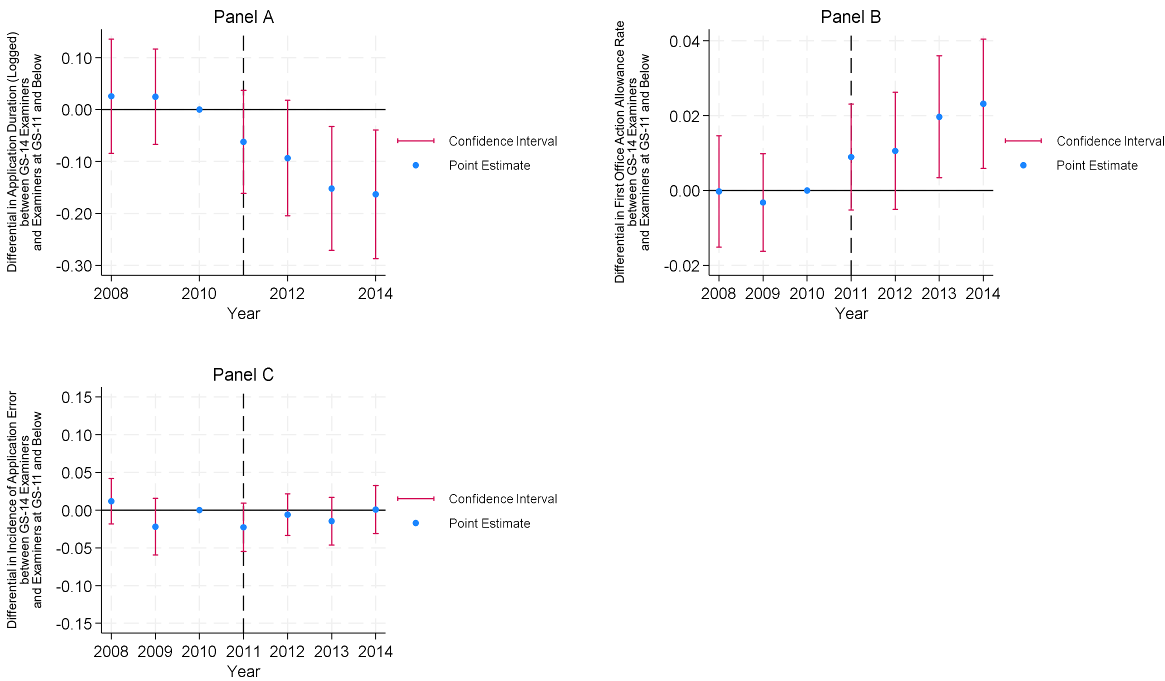

Notes: this figure replicates Figure 7 in the text except that it controls additional application-level controls for the number of dependent claims, number of independent claims, number of words in dependent claims, number of words in independent claims, average number of words in dependent claims, average number of words in independent claims, minimum number of words per claim among dependent claims and minimum number of words per claim among independent claims.

Overall Reform-Effects Analysis: Balance in Application Characteristics

One potential concern with the difference-in-difference analysis presented in Part V of the paper is that the nature of the applications assigned to GS-14 examiners (treatment examiners) may be diverging around the time of the reform relative to the nature of the applications assigned to GS-11-and-below examiners (control examiners), which could explain the observed increases in first-round allowance rates for treatment relative to control examiners after the reform and associated reductions in examination pendencies. To begin, of course, there is no reason to expect as much given that the application assignment process within Art Units (which are controlled for) are tangential to the patent worthiness of the application, as discussed in Part II. Nonetheless, to further appease these concerns, we demonstrate that the effects of the reform presented in Part V cannot likely be explained by changing application characteristics.

To demonstrate this, we focus on the first-round review allowance rate results. Further, we attempt to demonstrate application-balance in an omnibus fashion by forming the predicted likelihood that a first-round review is allowed based on observable application characteristics. In particular, we regress the incidence of a first-round allowance on an indicator for the applicant being a small entity and on a series of claim characteristics (and their squares): number of independent claims, number of dependent claims, total number of words in independent claims, total number of words in dependent claims, average number of words in independent claims, average number of words in dependent claims, minimum word count among word counts in independent claims, and minimum word count among word count in dependent claims. We then estimate the difference-in-difference specification set forth in Part V using this predicted first-round allowance rate as the dependent variable (without including these characteristics as right-hand-side controls). We present the results of this exercise in Panel A of Figure A22 below. In Panel B, we replicate the results from Figure 7 in the text, which explores the differential impact of the 2011 reform for treatment relative to control examiners on the *actual* incidence of a first-round allowance.

As demonstrated, while we find a notable increase in first-round allowance rates after the reform for treatment relative to control examiners, we observe no difference at all in predicted rates of first-round allowances. Accordingly, the pattern depicted in Figure 7 is not likely driven by changing observable application characteristics. This finding, combined with details on the assignment process lends confidence to the key identifying assumption that the first-round review outcomes for the treatment and control examiners would have otherwise trended in the same direction but for the reform.

FIGURE A22: EFFECT OF 2011 REFORM ON GS-14 RELATIVE GS-11-AND-BELOW EXAMINERS, SEPARATELY FOR PREDICTED FIRST-ROUND ALLOWANCE RATES (BASED ON APPLICATION CHARACTERISTICS) AND ACTUAL FIRST-ROUND ALLOWANCE RATES

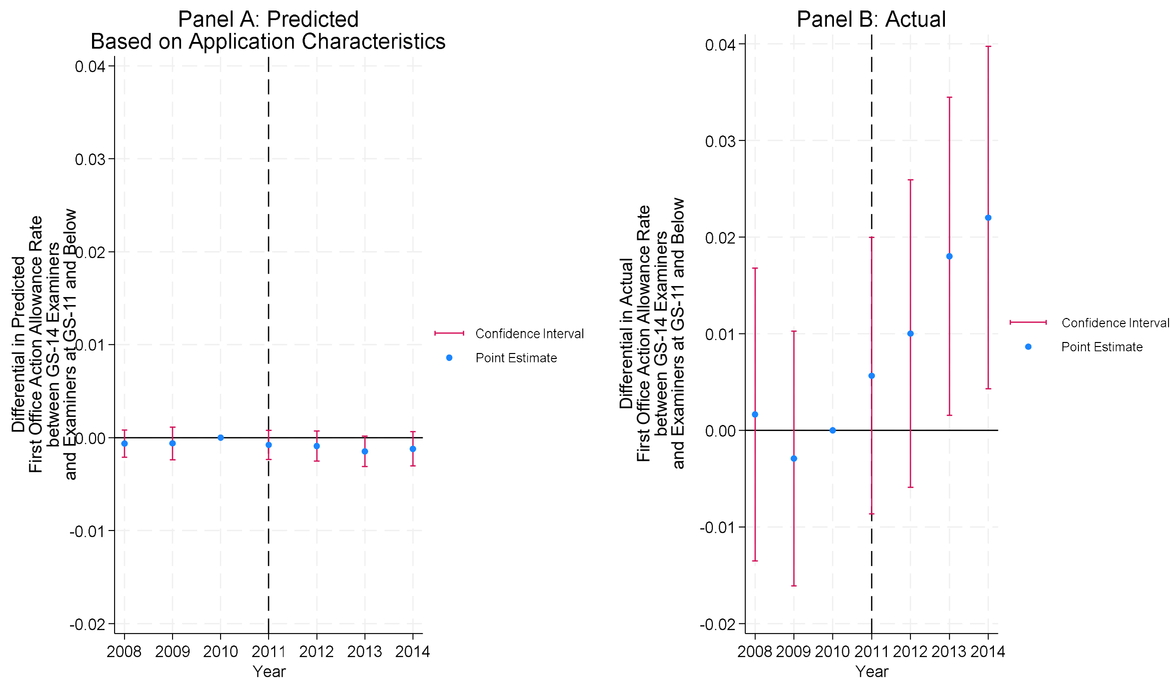

Notes: Panel B of this figure replicates Panel B of Figure 7 of the text. Panel A estimates a specification similar to that underlying Panel B of Figure 7 except that it uses the predicted incidence of a first-round allowance where this prediction is based on a range of application characteristics (e.g., claim counts and lengths). N = 1,667,802 in Panel B and 1,543,036 in Panel A (smaller than Panel B because we don't have claim characteristics for all applications).

Additional Mechanism Analyses: Triple-Differences Analyses

FIGURE A23: DIFFERENTIAL EFFECT OF 2011 REFORM ON EXAMINATION DURATION AND FIRST-ROUND ALLOWANCE RATES FOR ART UNITS WITH HIGH-VERSUS-LOW PRE-REFORM END-LOADING TENDENCIES AND FOR GS-14 RELATIVE TO GS-11-AND-BELOW EXAMINERS

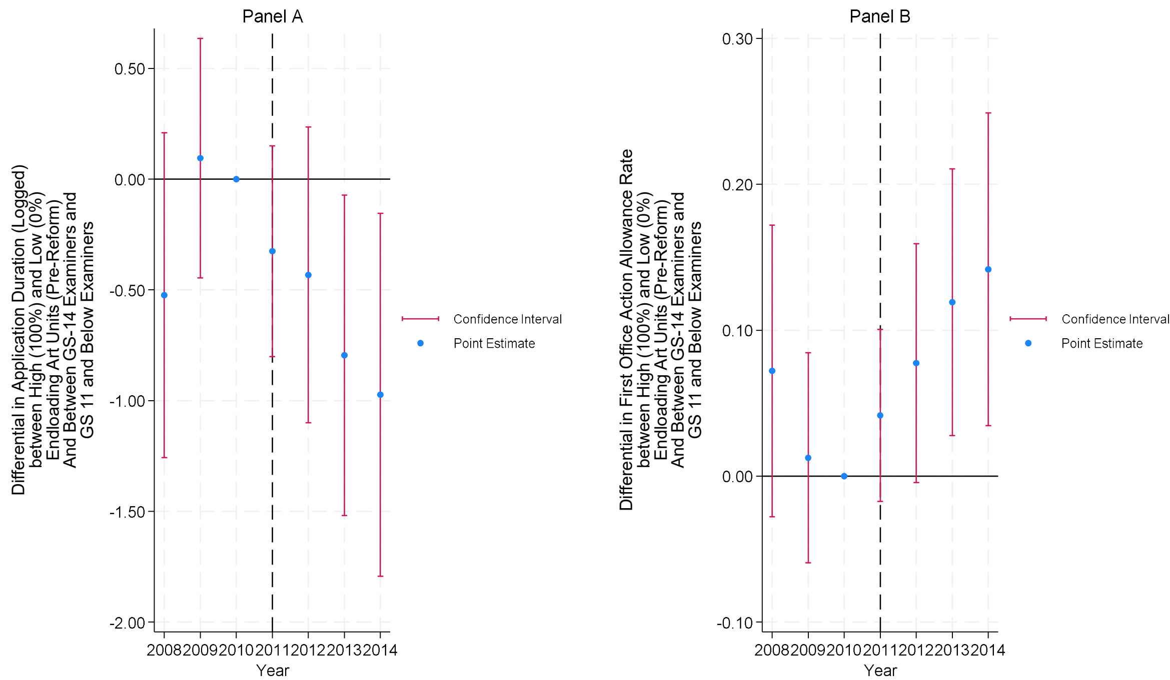

Notes: this figure shows results of a triple-differences specification, comparing the indicated outcome year-by-year before and after the 2011 reform, separately by examiner treatment status (GS-14 versus GS-11-and below) and separately by the pre-reform end-loading level of the examiner's art unit (measure continuously between 0 and 1). The reported coefficients are of the triple interaction terms, leaving 2010 out as the reference period. The underlying specification includes the various constitutive terms of the triple interaction along with the other controls included in specification (2) of the text. Standard errors are clustered at the Art Unit level.

Examiner Retention

Figure A24: Number of Actively Reviewing Examiners by Treatment Status in Years Leading up to and Following 2011 Patent Office Reform