A Report on "The response of deferred executive compensation to changes in tax rates" by Gorry et al. (2017)

Reviewer 2

February 22, 2026

v1

isitcredible.com

Disclaimer

This report was generated by large language models, overseen by a human editor. It represents the honest opinion of The Catalogue of Errors Ltd, but its accuracy should be verified by a qualified expert. Comments can be made here. Any errors in the report will be corrected in future revisions.

> I am wiser than this person; for it is likely that neither of us knows anything fine and good, but he thinks he knows something when he does not know it, whereas I, just as I do not know, do not think I know, either. I seem, then, to be wiser than him in this small way, at least: that what I do not know, I do not think I know, either.
>
> Plato, *The Apology of Socrates*, 21d

> To err is human. All human knowledge is fallible and therefore uncertain. It follows that we must distinguish sharply between truth and certainty. That to err is human means not only that we must constantly struggle against error, but also that, even when we have taken the greatest care, we cannot be completely certain that we have not made a mistake.
>
> Karl Popper, 'Knowledge and the Shaping of Reality'

Overview

Citation: Gorry, A., Hassett, K. A., Hubbard, R. G., & Mathur, A. (2017). The response of deferred executive compensation to changes in tax rates. *Journal of Public Economics*, Vol. 151, pp. 28–40.

Abstract Summary: This paper examines how taxes influence the choice of compensation and documents that income deferral is an important margin of adjustment in response to tax rate changes, particularly given the increasing use of stock options in executive compensation. The empirical analysis explores deferral by estimating how executives' choice of compensation between current and deferred income depends on changes in tax policy, suggesting a significant impact of taxes on the composition of executive compensation.

Key Methodology: The paper uses data from Execucomp for the period 1992-2005 and employs a two-stage least squares approach with executive-firm fixed effects to estimate how the share of deferred income responds to changes in the log of the tax price, using the permanent income tax rate as an instrument.

Research Question: How do changes in tax rates influence the choice of executive compensation, specifically the deferral of income?

Summary

Is It Credible?

The article examines how changes in tax rates influence the structure of executive compensation, arguing that high-income earners actively adjust the composition of their pay to minimize tax burdens. The authors claim that "income deferral is an important margin of adjustment in response to tax rate changes" (p. 28). Specifically, they find that higher personal tax rates lead to a significantly greater share of stock options in total compensation, while restricted stock grants decrease. The study also highlights the impact of the 1993 Section 162(m) rule—which capped the corporate deductibility of non-performance pay at $1 million—arguing that it provided a strong incentive for firms to shift compensation toward exempt stock options. Finally, the authors construct a dynamic model to illustrate that this deferral behavior generates efficiency costs and revenue losses, suggesting that the static elasticity of taxable income is insufficient for welfare analysis.

While the article identifies compelling correlations and offers a sound theoretical mechanism, the causal interpretation of its empirical estimates faces several significant methodological hurdles. The central identification strategy relies on using an executive's "permanent income tax rate" as an instrument for their current marginal tax rate. For this instrument to be valid, it must satisfy the exclusion restriction, meaning permanent income can only affect the composition of compensation through its impact on the current tax rate. This is a strong assumption. Executives with higher permanent income likely possess greater bargaining power, different risk preferences, or work for firms with distinct compensation philosophies that inherently favor equity-based pay. If these unobserved, time-invariant characteristics influence the share of options, the exclusion restriction is violated, and the estimates may be biased. While the authors attempt to address this by using an annual in-

come instrument in robustness checks, the substantial difference in point estimates between the two methods suggests the results are sensitive to the identification strategy.

The study's vulnerability to confounding historical events is another major limitation. The sample period of 1992–2005 coincides with the dot-com bubble and a massive, non-tax-related corporate governance movement that championed "pay-for-performance" and shareholder value maximization. To preserve variation in the federal tax rate, the authors omit year fixed effects in favor of a linear time trend (p. 33). A linear trend cannot adequately capture the profound, non-linear shocks of the 1990s stock market boom or sudden, firm-specific shifts in governance norms. While the authors include controls for firm market value and stock returns, and test specifications with future tax rates to capture anticipation effects, the model still risks incorrectly attributing the dramatic rise in option-based pay to tax policy changes rather than to these powerful concurrent market and institutional forces.

Furthermore, the variable used to capture the effect of the $1 million deductibility rule introduces potential mechanical and endogeneity issues. The variable is constructed as the share of total compensation exceeding one million dollars (taking a value of zero if compensation is below the threshold). Because total compensation serves as the denominator for both this independent variable and the dependent variable (the share of deferred income), any unobserved shocks or measurement errors in total compensation will affect both, potentially inducing a spurious correlation. Additionally, this variable inherently acts as a proxy for being a highly paid executive, a status that naturally correlates with a higher proportion of equity-based incentive pay regardless of tax considerations. While the authors argue this variable captures the specific deductibility channel—supported by its significance for deductible options but not for non-deductible restricted stock—the mechanical relationship remains a concern.

Finally, there is a tension between the article's dynamic theoretical framework and

its primary empirical approach. The authors motivate the study with the "option value" of deferring income under tax uncertainty, but their main regression models are largely static, relating current compensation structure to current tax rates. The authors do extend the analysis to include future tax rates (Table 5), finding evidence of anticipation, which supports their dynamic story. However, the theoretical model in the appendix, which elegantly illustrates the welfare implications, relies on parameters calibrated over a very short time horizon (1992–2007) and an elasticity of intertemporal substitution explicitly chosen to ensure the model's behavior is "consistent with what is observed in the data" (p. 38). Ultimately, the article provides valuable descriptive evidence of how executive compensation evolved during a period of shifting tax rules, but the precise magnitude of the causal behavioral response remains uncertain due to the complex macroeconomic and institutional changes occurring simultaneously.

The Bottom Line

The article argues that higher personal tax rates cause executives to defer income by shifting their compensation toward stock options. While the theoretical mechanism is intuitive and the descriptive trends are clear, the causal estimates are highly vulnerable to confounding factors. The reliance on a linear time trend during the dot-com bubble and the 1990s corporate governance revolution makes it difficult to isolate the effect of tax policy from broader market forces. Consequently, while taxes undoubtedly play a role in compensation design, the precise structural magnitude of the behavioral response estimated in this study should be interpreted with caution.

Potential Issues

Validity of the instrumental variable for personal tax rates: The article's central claim to causality relies on using the "permanent income tax rate" as an instrument for the current marginal tax rate, but the validity of this instrument is debatable. The instrument is constructed using an executive's average total compensation for all years they appear in the database (p. 32, footnote 15). For the instrument to be valid, it must satisfy the exclusion restriction, meaning it can only affect the share of deferred compensation through its effect on the current tax rate. This assumption may be violated if an executive's permanent income level is correlated with unobserved, time-invariant factors that also determine compensation structure. For instance, executives with higher permanent income may possess greater bargaining power, different risk preferences, or work for firms with specific compensation philosophies that favor equity-based pay, independent of the current tax rate. These unobserved characteristics would reside in the error term and be directly correlated with the instrument. The authors do address this by running robustness checks using the annual marginal tax rate as an instrument, but the potential bias from a violated exclusion restriction in the primary specification remains a concern (p. 35, Table 4 notes).

Omission of year fixed effects and confounding historical events: The study's identification strategy may be vulnerable to confounding from major historical events due to the omission of year fixed effects. The authors state they use a linear time trend instead of year dummies because including the latter would absorb the limited time-series variation in the federal tax rate, making the effect of tax policy changes difficult to estimate (p. 33, footnote 17). While this is a common econometric trade-off, it leaves the model susceptible to bias from any aggregate shocks that do not follow a simple linear trend. The sample period of 1992–2005 coincides with the dot-com stock market bubble and its subsequent collapse, a profound, non-linear event

that dramatically altered the use and valuation of stock options for reasons unrelated to personal tax policy. The model includes controls for firm market value and assets, which may capture some of these shocks, and the authors introduce a specification with future tax rates to test for dynamic effects (p. 35). However, a simple time trend and firm-level controls are unlikely to fully account for such powerful aggregate shocks in the way that a full set of year fixed effects would. This creates a risk that the model incorrectly attributes changes in compensation structure caused by the stock market to the tax policy changes that occurred around the same time.

Limited control for corporate governance trends: The analysis may not adequately account for the powerful, non-tax-related influence of the corporate governance movement during the sample period. The 1990s and early 2000s saw a strong push for "pay-for-performance" and "shareholder value maximization," which strongly favored stock options as a tool to align executive and shareholder interests. This provides a significant alternative explanation for the observed rise in option-based pay. The article acknowledges this competing explanation and attempts to control for it using firm-executive fixed effects and a linear time trend (p. 29). However, these are blunt instruments. Firm fixed effects capture time-invariant governance philosophies, and a linear trend captures a smooth, secular shift. Neither can account for time-varying, firm-specific changes in governance, such as a board facing sudden pressure from an activist investor to increase the share of equity-based pay. Without more direct proxies for governance quality or shareholder pressure, it is difficult to disentangle the effect of tax policy from these concurrent shifts in governance norms.

Potential endogeneity of the "million-dollar rule" variable: The variable used to measure the effect of the Section 162(m) "million-dollar rule" may be endogenous. The variable is constructed as the share of total compensation exceeding one million dollars (taking a value of zero if compensation is below the threshold) (p. 32). The dependent variable is the share of deferred income in total compensation. Because

the independent variable is a non-linear function of the denominator of the dependent variable, any unobserved shocks or measurement error in total compensation will affect both variables, potentially inducing a spurious correlation. Furthermore, the variable is a proxy for being a very highly paid executive. It is a well-established pattern that as executive pay increases, the proportion of equity-based compensation also increases for incentive-alignment reasons, independent of tax considerations. While the authors argue that the variable captures the specific deductibility channel—noting its significance for deductible options versus non-deductible restricted stock—the model's controls may not fully capture the non-linear relationship between the level of pay and its composition, meaning the strong coefficient on this variable could reflect this confounding effect rather than a pure causal effect of the tax rule.

Sensitivity of the key coefficient across specifications: The magnitude of the article's main coefficient of interest—the effect of the tax price on the share of stock options—appears sensitive to the choice of instrumental variable, raising questions about the robustness of the point estimate. In the baseline specification using the permanent income instrument, the coefficient is -0.403 with a standard error of 0.202 (Table 4, p. 35). In a robustness check using an annual total income instrument, the coefficient's magnitude increases to -0.705 with a standard error of 0.319 (Table 4, p. 35). While the sign and statistical significance of the effect are consistent across both models, the point estimates differ substantially. However, given the large standard errors, the confidence intervals of the two estimates overlap significantly, suggesting the difference may not be statistically meaningful. Nonetheless, the sensitivity of the point estimate suggests that the precise quantitative estimate of the behavioral response may depend on the specific identifying assumptions of the instrument, warranting caution in interpreting its exact structural magnitude.

Mismatch between dynamic theory and the primary empirical model: There is a potential disconnect between the article's theoretical motivation and its primary

empirical specification. The analysis is motivated by the dynamic "option value" of deferring income under tax uncertainty, a concept explored in a formal model in Appendix B (p. 38). This framework implies that an executive's decision depends on the entire expected future path of tax rates. However, the main empirical model presented in Table 4 is largely static, relating the current compensation structure to the current tax rate (p. 35). The authors do extend the model in Table 5 to include leads of the tax rate for one and two years in the future, explicitly stating that the purpose is to test whether tax effects are reversed or persist over time (p. 36). While this specification directly tests for dynamic anticipation and reversal, it remains a limited, reduced-form approximation of the complex dynamic optimization problem described in the theory. As a result, the conclusion that the findings represent a "long-term change in the nature of compensation" rather than a "simple timing shift" may be an overstatement based on this limited dynamic specification (p. 36).

Interpretation of the corporate tax rate finding: The article's conclusion that there is "little evidence that the corporate tax rate influences deferral decisions" is presented as a finding consistent with the model's specification (p. 29). The regressions include the "Net Income Share over $1 Million" variable, which is designed to capture the primary channel through which corporate taxes affect executive compensation: the deductibility limit imposed by Section 162(m). The authors explicitly argue that once this powerful deductibility channel is accounted for, theory suggests the general corporate tax rate should have less influence (pp. 29, Appendix A). However, the model also includes a firm-specific marginal tax rate variable, α, that accounts for tax loss carryforwards (p. 34). The lack of significance on α suggests that even for firms with different effective tax rates (due to losses), the decision to defer is not driven by the corporate rate *per se*, but by the specific deductibility rule. Therefore, the finding is less a novel discovery and more a confirmation that the main corporate tax effect operates through the specific channel of the deductibility cap.

Justification for dynamic model parameters in the appendix: The dynamic model

presented in Appendix B relies on calibrated parameters whose justification could be stronger. The stochastic process for the tax rate is derived from an AR(1) process estimated using statutory top marginal rates from 1992 to 2007, which is a very short time series from which to reliably estimate a persistent macroeconomic process (p. 38). Furthermore, the choice for the elasticity of intertemporal substitution (EIS) is justified by stating that an elasticity greater than one is chosen so that deferral behavior is "consistent with what is observed in the data" (p. 38, footnote 26). While calibrating parameters to match observed behavior is a standard practice for simulation models, this justification clarifies that the parameter regime was selected to ensure the model produced the desired qualitative result, rather than being drawn from external evidence. The authors do acknowledge that their chosen value of 2 is relatively high compared to some estimates in the literature.

Presentation and transparency issues: The article has several minor issues related to presentation and transparency. First, the analysis does not report the results of a simple OLS regression of deferred compensation on the uninstrumented current tax rate. While OLS is known to be biased, presenting it is standard practice as it allows readers to assess the direction and magnitude of the endogeneity bias the instrumental variable is intended to correct; the authors note a raw correlation in a footnote but do not present the regression (p. 34, footnote 20). Second, the dataset on executive compensation is known to be highly skewed. The authors note that they exclude a few famous outliers like Steve Jobs, but the article does not report a systematic procedure for handling influential observations, such as winsorizing or trimming the data at a specified percentile, which is common practice in this literature (p. 33, footnote 18). The use of robust standard errors clustered by year and firm-executive pair does, however, address some distributional concerns (p. 32).

Future Research

Alternative identification strategies: Future research could exploit state-level tax variations or specific, staggered tax reforms that allow for the inclusion of a full set of year fixed effects. This would help absorb aggregate macroeconomic shocks, such as stock market bubbles, and secular trends in corporate governance, thereby isolating the specific causal impact of tax policy on compensation structure.

Refined measures of regulatory impact: To better understand the impact of corporate deductibility limits like Section 162(m), researchers could utilize regression discontinuity designs around the $1 million threshold. Alternatively, studies could explore firm-level variations in the binding nature of the rule using instruments that do not mechanically rely on the denominator of the dependent variable, reducing the risk of spurious correlation.

Dynamic structural estimation: Future work could estimate a fully dynamic structural model of executive compensation that explicitly incorporates executives' expectations of future tax rates, stock market returns, and vesting schedules. Moving beyond reduced-form approximations would allow researchers to directly quantify the "option value" of deferral and provide more precise estimates of the dynamic welfare costs associated with tax uncertainty.

isitcredible.com