EJW
ECON JOURNAL WATCH
Scholarly Comments on
Academic Economics

ECON JOURNAL WATCH 20(1)
March 2023: 34–44

# Response to Gamino

Scott Barkowski[1] and Joanne Song McLaughlin[2]

LINK TO ABSTRACT

Aaron Gamino (2023) comments on our paper analyzing marriage of young adults and dependent coverage health insurance mandates in the *Journal of Human Resources* (Barkowski and McLaughlin 2022), arguing that we used "potentially misspecified difference-in-difference-in-differences models" that "omit relevant interaction terms" (Gamino 2023, 17). In particular, he argues our model is biased by the exclusion of age-by-year interaction terms. When he includes these terms in our model, he obtains coefficient estimates that are close to zero for some model variations. In this reply, we argue that Gamino's inclusion of age-by-year interaction terms in our model is inappropriate and introduces bias, and that his analysis generally *supports* the findings in our paper.

To summarize our arguments, there are at least four important flaws underlying Gamino's critique.[3] First, he ignores our discussion (Barkowski and McLaughlin 2022, 641–643, 648 Figure 1) of how there is no true control group that is not affected by the treatment in our data, and how this changes our analysis from a standard difference-in-difference (DD) or triple-difference (DDD) research design. His introduction of age-by-year fixed effects to our model, consistent with treating our research design as a DDD, introduces bias to one of the main coefficient estimators (and hence, the entire model). Second, rather than removing bias, Gamino's additional fixed effects absorb the remaining identifying variation in the model. Evidence for this comes from his analysis of our marriage entry variable. Our model explains much less of the variation in marriage entry

1. Clemson University, Clemson, SC 29634.
2. University at Buffalo, Buffalo, NY 14260.
3. We were given extremely limited time to respond concurrently to Gamino's comment. Hence, our discussion and detail provided in our arguments is necessarily limited. Additionally, given the limited time, we have not reviewed or checked for errors any of Gamino's Stata code or confirmed his results at all.

than our standard outcome, marriage state. When he adds age-by-year fixed effects to the marriage entry model, estimates are either little changed or become larger in magnitude. Third, Gamino performs a placebo policy analysis to estimate bias due to omitted age-by-year fixed effects. The bias he finds suggests our estimates may understate the magnitude of the effects. This is inconsistent with the true effects being close to zero, as he argues, and represents additional evidence that his added fixed effects merely absorb important variation. Fourth, Gamino provides minimal justification for why he thinks our model is inappropriate. Gamino (2023, 18) implies we use a DDD model in a corresponding analysis in a related paper. This claim is incorrect, as the analogous model in the related paper (Barkowski, McLaughlin, and Ray 2020) is a DD model, not a DDD model. Lastly, while Gamino argues our analysis of marriage with a DD-style model is inappropriate, in his own work studying dependent coverage under the same state-level mandates (Gamino 2018), he uses a DD model when he considers marriage, and also does not include all possible multi-way-interaction terms in his analyses of other outcomes. Clearly, Gamino agrees that the type of model used in an analysis should be based on the context, and that DD-style models are appropriate in the current one, even if it does not include all possible multi-way-interaction fixed effects.

In the rest of this reply, we provide more in-depth discussion of the above points. However, we also urge readers to consult our original paper, where we provide an extensive discussion on the rationale behind our model and its identification. Clearly, every model has weaknesses and limitations, and as ours does not include age-by-year interaction terms, it could be affected by unobserved age-year trends—a point that was acknowledged in our paper (Barkowski and McLaughlin 2022, 684, 684 n.29). But our model was not "misspecified" or lacking the "proper" controls (Gamino 2023, 16, 18). Rather, as Gamino should have been aware, it was chosen for specific reasons that were fully discussed.

## Bias is introduced by the addition of age-by-year interactions

Our analysis of marriage focuses on the interaction of federal and state-level policies that mandated increased ages at which young adults could be enrolled on their parents' health insurance plans. The state-level mandates came first and were implemented at various points across states starting in the 1970s, while the federal mandate came into being in 2010 as part of the Affordable Care Act (ACA). The interaction of these two types of policies is of interest because the state-level policies often limited eligibility on the basis of factors like marital status and college enrollment, whereas the ACA policy did not. In the case of marriage, under state-level policies, a young adult was required to be unmarried in order to be eligible to be enrolled on his or her parents' coverage. This created a disincentive to be married for young adults. In contrast, the ACA mandate allowed young adults to be on their parents' coverage regardless of their marital state. For young adults eligible for coverage under pre-existing state-level mandates, the ACA's enactment effectively ended the prohibition against marriage, resulting in an incentive, relatively speaking, towards marriage. For others not affected by state mandates, though, the ACA created a disincentive to marry because it opened parents as a potential source to obtain health insurance coverage. For young adults that might have considered marrying to obtain health insurance coverage (through, for example, a spouse's employer) this parental source offered an easier alternative. Hence, for those not eligible under state mandates, the ACA disincentivized marriage (Abramowitz 2016).

Our analysis of these policies was designed to identify these contrasting effects via a model we called a "DD style" model (Barkowski and McLaughlin 2022, 644, 645). We start by noting that eligibility for state-level mandates depends primarily on three factors: the state of residence, age, and year. We use only these factors to determine eligibility in our analysis because additional factors required in eligibility, like college enrollment, are endogenous and one of the other primary factors, marriage, is our main outcome variable of interest. Defining eligibility this way, we then used a model of the following form:

$$Y_{iast} = \beta_1 ELIG_{ast} \times ACA_t + \beta_2 ELIG_{ast} + X'_{iast}\gamma + a_a + \delta_{st} + U_{iast}\,.$$

We estimate this model primarily on data for 19- to 25-year-olds using American Community Survey and Census data for 2000 to 2015. Note that unlike most analyses of the ACA dependent coverage mandate, we do not use data on individuals older than 25 in our main estimates.

We called the equation above a "DD style" model and not a DD model because, strictly speaking, it is not a DD model since there is no unaffected control group included in our analysis. Everyone in our sample is treated by the ACA mandate (since no one over 25 is included in our primary sample). This would typically mean that this model would not be appropriate in standard causal effect analysis. In our paper, however, we demonstrate that we can still use it to identify the interaction effect of the ACA and state-level dependent-coverage mandates (see Barkowski and McLaughlin 2022, 641–643, 648 Figure 1).

In that demonstration, we note that $\beta_2$ measures the marriage gap before the implementation of the ACA between those who were eligible for state mandated coverage and those who were not, while $\beta_1 + \beta_2$ measures the same after the ACA.

We should expect $\beta_2$ to be negative because state mandates discouraged marriage, but $\beta_1$ should be positive because the ACA simultaneously encouraged marriage for those affected by state mandates and discouraged it for those who were not. That is, $\beta_1$ should be positive because we expect that the ACA narrowed the marriage gap, and $\beta_1$ measures the difference between the post-ACA and pre-ACA marriage gaps. See Figure 1 in our paper for a stylized illustration of this argument (Barkowski and McLaughlin 2022, 648).

Table 1 below presents a summary of how estimates of $\beta_1$ are affected by using a DD-style or DDD-style model. The table shows expected effects on marriage by eligibility groups: whether an individual is the older or younger group (eligible individuals were younger) and whether a person is in a mandate or non-mandate state. We call the effect on each respective group A, B, C, and D, and within the table give a brief prediction of the effect and reason for it. Note that these focus on the interaction effect of the ACA, so they represent the change from before the ACA implementation to after. Thus, the effects are changes, meaning they already incorporate the first difference of a DD- or DDD-style model.

**TABLE 1. Summary of predicted interaction effects on marriage for eligibility groups**

| | **State type** | |
|---|---|---|
| **Individual age** | Mandate | Non-mandate |
| Younger (Eligible) | Effect **A**: Increase (Marriage prohibition removed) | Effect **C**: Decrease (Parents better health insurance source than marriage) |
| Older (Ineligible) | Effect **B**: Decrease (Parents better health insurance source than marriage) | Effect **D**: Decrease (Parents better health insurance source than marriage) |
| **Interpretation of estimators in DD- and DDD-style models** | | |
| In a DD-style model: | $\hat{\beta}_1$ ≈ Effect **A** − (average of Effects **B**, **C**, and **D**)<br>≈ Positive − Negative<br>≈ Positive + Positive<br>≈ **Positive Overall** | |
| In a DDD-style model: | $\hat{\beta}_1$ ≈ Effect **A** − Effect **B** − (Effect **C** − Effect **D**)<br>≈ Positive − Negative − (Negative − Negative)<br>≈ Positive + Positive − **(Negative + Positive)**<br>≈ DD-style model effect − **(Negative + Positive)**<br>≈ Total effect will deviate from effect estimated in DD-style model if Effects **C** and **D** do not cancel out. | |

In the bottom part of the table we summarize the effects estimated by the different DD- or DDD-style models. A DD-style model gives an estimate approximately equal to Effect **A** minus an average of Effects **B**, **C**, and **D**. Because the predicted effects **B**, **C**, and **D** are all negative, this identifies the interaction effect of interest, the effect on those with state mandates plus the effect on those without them. When state-by-year interactions are included, the estimate is given by Effect **A** − Effect **B** because only within-state variation is used, but the same rationale follows.

When a full DDD-style model is used, the estimator uses the mandate and non-mandate states to compute separate DD-style estimates, and then differences them. Hence, the estimate is approximately given as Effect **A** − Effect **B** − (Effect **C** − Effect **D**). The first part of this, Effect **A** − Effect **B**, represents an estimate of the interaction effect of interest, but the DDD-style estimate includes Effect **C** − Effect **D**, which means the DDD-style estimate does not estimate the effect of interest unless Effects **C** and **D** have the same magnitude and direction, and, hence, cancel out. Thus, if older and younger individuals in non-mandate states have different effects from the implementation of the ACA, the estimator for $\beta_1$ Gamino uses in his comment will be biased as an estimator of the interaction effect of interest.

There are clear reasons to think that Effects **C** and **D** are likely different. Older individuals are more likely to have potential spouses with whom they could alter their marriage plans to access health insurance. This suggests the alternative access granted by the ACA is more likely to affect marriage among older individuals. In contrast, younger individuals are less likely to have employers that offer employee health insurance coverage, so parental coverage represents the first alternative to marriage for younger people in this situation to access group health insurance. This suggests marriage is more likely to be affected for younger people. Our work on the mandates' effects on health insurance coverage found that younger individuals were more likely to have their insurance coverage affected by the state mandates (Barkowski, McLaughlin, and Ray 2020). If younger individuals were also more likely to take up coverage under the ACA mandate, we would again expect a bigger effect for younger individuals.

In sum, the ACA's effects on older and younger individuals in non-mandate states are likely to be heterogeneous, and Gamino's estimator is, therefore, biased. Further, since we do not have a good prediction for which group effect is likely to be larger, the bias could go in either direction. A proper analysis should use DD-style models like those used in our paper.

## The introduction of age-by-year fixed effects absorbs the remaining identifying variation

In Gamino's replication of our analysis, he adds age-by-year fixed effects to two slightly different models of marriage. The first is the model used for our main results, where marital state is used as the outcome. That is, whether the respondent was married at the time of the survey interview. The second model uses entry into marriage—that is, whether the person was married in the last year—as the outcome. For both of these outcome variables, our analysis uses the same model with the same controls, and the theoretical predictions are the same in both cases. However, our model explains far less of the variation in marriage entry than in marital state. This makes sense given that, in one sense, marriage entry is a change or first-differenced version of marital state. As a result, the variation explained in our marriage entry outcome by our DD-style models is only about 15 percent as much as the variation explained in marital state.

We mention the above because one of the persistent concerns with using fixed effects is that they can end up absorbing all of the good identifying variation in a research design. In the case of Gamino's comment, this seems to be the case because when he adds age-by-year interactions to our model of marital state, he obtains estimates that are small and close to zero (see Gamino 2023, 19 Table 1)—consistent with what happens when one absorbs all the identifying variation. In contrast, when he adds the same fixed effects to the model of marriage entry, he obtains results similar to our original estimates. This suggests that in the model absorbing much more variation, the additional fixed effects are wiping out the identifying variation, but not in the model that explains less of the variation.

To show the similarity between our marriage entry results with and without age-by-year fixed effects, Table 2 below presents Gamino's replication of our original results and his estimates after adding the additional controls. It includes, in columns 3 and 4, results from an expanded model that was included in our original analysis for more direct comparison with Joelle Abramowitz's (2016) results. It included individuals ages 27–30 so that a DD comparison could be made between individuals affected by the ACA (below 26) and those who were not (above age 25). The estimates under the "ACA DD" heading are based on this comparison. The estimates beneath the "Under age 26" measure the marriage gap between the state eligible and ineligible who were below 26. That is, they have the same interpretation as the main coefficients of interest in our primary model.

Notably, our result from the expanded version of our model (columns 3 and 4), is more consistent with our theoretical predictions when the age-by-year fixed effects are included. This is seen most clearly under the ACA DD heading, where the estimate for the state eligible moves from being negative to positive. Additionally, the "Under age 26" estimates are similar with and without the age-by-year fixed effect. In his comment, Gamino focuses on the estimates in column 2 in our Table 2 being statistically insignificant, but this is a shortsighted interpretation. Gamino is claiming that omitting age-by-year interactions is causing bias, but his work here shows little evidence of such bias. Moreover, the fact that the estimates in column 3 are a *little bit* smaller and less precise is not surprising given that the added fixed effects absorb a large amount of variation.

**TABLE 2. Gamino's addition of fixed effects gives results similar to our original results for the marriage entry outcome**

| Outcome = Marriage entry | Original model and sample: Ages 19–25 | | Expanded model and sample: Ages 23–25 and 27–30 | |
|---|---|---|---|---|
| | Replication of original result | With Gamino's added age-year FE | Replication of original result | With Gamino's added age-year FE |
| | (1) | (2) | (3) | (4) |
| Under age 26 | | | | |
| Eligible under state mandates × ACA **Prediction: (+)** | 0.0049* (0.0029) | 0.0042 (0.0029) | 0.0056** (0.0022) | 0.0075*** (0.0024) |
| Eligible under state mandates **Prediction: (−)** | −0.0074** (0.0031) | −0.0052 (0.0040) | −0.0071** (0.0033) | −0.0071** (0.0035) |
| ACA DD | | | | |
| Effect on state ineligible **Prediction: (−)** | Not included in model | Not included in model | −0.0090*** (0.0017) | Not identified because of age-year FE |
| Effect on state eligible **Prediction: (+)** | Not included in model | Not included in model | −0.0029 (0.0030) | 0.0053 (0.0036) |
| FE included | state-year, age-state | state-year, age-state, age-year | state-year, age-state | state-year, age-state, age-year |
| Gamino (2023) table | Table 2, Col. 5 | Table 2, Col. 6 | Table 3, Col. 5 | Table 3, Col. 6 |

*Notes*: All results in this table come from Gamino's replication of our original analysis.

In our view, the key takeaway from Gamino's exercise of adding fixed effects to our models and obtaining contrasting results when marital state and marriage entry are used as the outcome variables is that the fixed effects are absorbing the identifying variation in one model, but not in the other. This is in addition to the introduction of bias by the new fixed effects, as well. So Gamino's suggestion has both a bias problem and a practical implementation problem.

The next section provides additional evidence regarding the fixed effects wiping out identifying variation, and it also suggests that (at least some of) our original estimates might be understated.

## Gamino's placebo policy analysis suggests our original conclusions are valid

In the appendix of his comment, Gamino performs a placebo policy analysis to try to estimate the bias in our models. The idea here is to randomly assign eligibility criteria and then re-estimate our model many times, producing the distribution of the estimators. Of most interest is his analysis of our model that included both state-by-year and age-by-state fixed effects. Since only the age-by-year interactions are missing here, any bias observed from this model would have to be coming from age-by-year trends. Section 4 of Gamino's Figure A1 shows these results (Gamino 2023, 29). Interestingly, for $\beta_2$ it suggests there is a bias, but it is positive (that is, the mean of the distribution is positive). This implies that our estimates of $\beta_2$—which are negative—in our model with age-by-state and state-by-year terms may understate the magnitude of the true coefficient.

In the case of $\beta_1$, Gamino's placebo analysis once again suggests there is a positive bias. Unlike in the case of $\beta_2$, this bias goes in the same direction as our coefficient estimate. However, we note that the bias is still less than our point estimate, so even if one were able to correct the estimator for bias, our estimate would still be positive and consistent with our theoretical prediction. Additionally, since $\beta_1$ is defined relative to $\beta_2$, if one were to correct for the bias in $\beta_2$, all else equal the estimator of $\beta_1$ would have to be larger (more positive) to account for the more negative estimate of $\beta_2$. Thus, absent the bias $\beta_2$, we would have expected a larger estimate for $\beta_1$, and, in turn, a larger estimate of $\beta_1$ net of the bias seen in Gamino's Figure A1. Bringing these results together, Gamino's placebo analysis suggests that our estimates for the pre-ACA marriage gap was too small, and that after the ACA the gap narrowed. These are the same conclusions we obtained from our original analysis. Hence, Gamino's critique provides additional evidence in support for our paper's hypothesis.

It is also worth contrasting these placebo results with the regression estimates Gamino obtains from adding age-by-year interactions to our model of marital state, and his placebo analysis of the model that also includes age-by-year fixed effects. In those regressions, he finds estimates close to zero, and in the placebo analysis he obtains distributions of estimates that are all very small and close to zero—very different from the placebo results for the other models, which are all much more widely dispersed. But if Gamino's placebo results for the model with only age-state and state-year fixed effects are valid, the models in our paper are not biased towards zero. So if the added age-by-year terms were merely absorbing bias, we should not obtain zero estimates or placebo distributions crowded around zero. So why does he obtain them? The answer is easy: the zero estimates and zero-neighborhood distributions are exactly the results we expect when one absorbs all the identifying information with added fixed effects.

## Gamino gives minimal justification for critiquing our model

Gamino provides minimal explanation for why he thinks our models are misspecified, and he never suggests any specific omitted factors that he is concerned could cause bias. Instead, he asserts that we are inconsistent in how we employ our models between our marriage paper and our related paper on insurance coverage (Barkowski, McLaughlin, and Ray 2020). Gamino's assertion is incorrect. Our marriage paper that he criticizes focuses on the interaction of state and federal policies (as should be clear from the paper's title). In our insurance coverage paper, the model most closely analogous to the model in our marriage paper is a DD model in which we omit age-by-year interactions. These models are analogous because they both consider the interaction of the state and federal policies and both are estimated using data spanning the introduction of the ACA dependent coverage mandate. A DDD model used in our insurance coverage paper is used only on data preceding the ACA and does not involve any interactions with the ACA mandate. In sum, Gamino's statements completely misconstrue the relationships between models in our papers.

Gamino further claims that our model is misspecified because the main variable of interest, *ELIG*, varies across three dimensions, "which requires the inclusion of three sets of two-way interaction terms" (2023, 17), one set of which our model does not include. This seems to assert the existence of a requirement that if a model's variable of interest has $n$ dimensions of variation, the model must include all possible ($n-1$)-way interaction terms. However, in his own work on the state dependent coverage mandates (Gamino 2018), Gamino's version of *ELIG* varied along 5 dimensions,[4] yet he did not include any 4- or 3-way interaction terms in his models of various insurance coverage and labor market outcomes.

Similarly, while Gamino's overall critique is essentially that he thinks we should not have analyzed marriage with a DD-style model, when he performed a simple analysis of marriage in his paper, it was based on a DD model. In this analysis, he used a version of *ELIG* that he only allowed to vary along the state and time dimensions, ignoring the age dimension for unstated reasons. Thus, in more than one facet, Gamino's own work in the same context has relied on similar techniques to those we use. This undermines his unsupported assertions that our model is "misspecified."

4. Gamino's version of *ELIG* for these analyses varied by student and marital status as well as the age, year, and state dimensions our *ELIG* varies over. We did not use student status because it is endogenous and marriage is our outcome.

# Additional points

## The placebo analysis is not useful for inference

We note that in some places, Gamino (2023) also claims his placebo analysis is a form of randomization inference, but we disagree. He subjects his randomization to a series of restrictions to reflect features of state-level policies, but true randomization would be far less restricted, allowing the years of implementation and ages to vary much more. Further, his randomization ignores the ACA treatment. A true randomization inference study would have to include the uncertainty in when and how the ACA was implemented since our primary question of interest studies the interaction of the policies. Thus, his distributions do not reflect all the sources of uncertainty in the assignment process. So, in our view, the placebo analysis is not useful for studying the inference in our study.

We also note that Gamino never provides any argument that the standard errors we used are not correct. So in addition to his exercise not being an appropriate randomization analysis, there is also no reason to think one is needed.

## Our coding of Pennsylvania was correct

In footnote 4, Gamino claims that we wrongly coded Pennsylvania's 2009 state mandate as being effective in 2010. Pennsylvania's law went into effect in September 2009, but our coding of it as effective in 2010 is consistent with the policy we applied to all of our laws: if the law was effective in the second half of the year, we treated it as being effective at the beginning of the next year. This is also why we treated the ACA mandate as going into effect in 2011. This policy was stated in our paper (Barkowski and McLaughlin 2022, 667).

# Conclusion

We acknowledge that there are limitations and pitfalls in all research, but we hope we demonstrated that our model was appropriately applied for the given setting. In this reply, we showed that Gamino's suggested change to our model is inappropriate because it introduces bias and the textbook DDD model is not appropriate to answer the main question of the paper. The zero estimates he obtains for models of marital state are inconsistent with estimates he obtains for the marriage entry outcome, and his placebo analysis does not suggest a bias towards zero from unobserved age-by-year trends. These facts suggest the zero results for marital state are most likely due to absorbing necessary identifying variation. Overall, while we appreciate his desire to investigate the robustness of our finding and attempt to contribute to this literature, we do not find his analysis convincing. Rather, it seems to add evidence consistent with our original hypothesis.

## References

**Abramowitz, Joelle**. 2016. Saying, "I Don't": The Effect of the Affordable Care Act Young Adult Provision on Marriage. *Journal of Human Resources* 51(4): 933–960.

**Barkowski, Scott, and Joanne Song McLaughlin**. 2022. In Sickness and in Health: Interaction Effects of State and Federal Health Insurance Coverage Mandates on Marriage of Young Adults. *Journal of Human Resources* 57(2): 637–688.

**Barkowski, Scott, Joanne Song McLaughlin, and Alex Ray**. 2020. A Reevaluation of the Effects of State and ACA Dependent Coverage Mandates on Health Insurance Coverage. *Journal of Policy Analysis and Management* 39(3): 629–663.

**Gamino, Aaron M**. 2018. New Evidence on the Effects of Dependent Coverage Mandates. Working paper. **Link**

**Gamino, Aaron M**. 2023. Health Insurance Mandates and the Marriage of Young Adults: A Comment on Barkowski and McLaughlin. *Econ Journal Watch* 20(1): 15–33. **Link**

## About the Authors

**Scott Barkowski** is assistant professor of economics at Clemson University. His email address is sbarkow@clemson.edu.

**Joanne Song McLaughlin** is associate professor of economics at the University at Buffalo. Her email address is jsmclaug@buffalo.edu.