

# Hazard Versus Linear Probability Difference-in-Differences Estimators for Demographic Processes

Lawrence L. Wu and Fangqi Wen

**ABSTRACT** This study examines the properties of the linear probability difference-in-differences estimator when the data are in fact generated by a single-decrement, continuous-time hazard process. We focus on the textbook case of two groups and two periods in which the control and treatment groups are observed before and after treatment. We provide formal derivations and illustrate matters concretely by reexamining economic studies that have relied on the linear probability difference-in-differences estimator when attempting to obtain estimates of the causal effect of unilateral and no-fault divorce. In particular, we show that the increasing then decreasing pattern of effects found by Wolfers (2006) can be generated by a time-invariant effect of treatment in a proportional hazard setting. We conclude that often implicit assumptions about how the data are generated are an important and necessary component of causal identification.

**KEYWORDS** Data-generating function • Difference-in-differences estimation • Dynamic response to treatment • Linear probability and proportional hazard regression • Unilateral and no-fault divorce

## Introduction

Difference-in-differences (DD) procedures are perhaps the method most heavily used to obtain plausibly causal estimates from observational data for treatments such as an exogenously imposed change in policy. The popularity of such DD procedures stems in no small part from their use with a wide range of data, including panels that follow individuals over time but also repeated cross sections for outcomes observed at the aggregate level for geographic units such as states or counties. Early examples include Ashenfelter and Card (1985) on the effect of job training programs on earnings and Card and Krueger (1994) on the effect of increases in the minimum wage on the demand for labor. Linear probability DD procedures have likewise been heavily used in analyses of binary outcomes.

In this study, we present formal results that question the use of the linear probability DD when the binary outcome of interest is a single-decrement hazard


process involving the transition from a common origin state to a single destination state. Examples include not only the traditional demographic outcomes of fertility, mortality, and migration, but also marriage, divorce, cohabitation, and spells of unemployment or program participation. For these and other hazard processes, life table methods and their regression extensions have been the method of choice in demography, epidemiology, public policy, sociology, and statistics, as well as in an older literature in economics that provided significant contributions to these methods.

We present formal derivations for the properties of the linear probability DD when the data are in fact generated by a continuous-time, single-decrement hazard process. We focus throughout on the textbook case of two groups and two periods in which the control and treatment groups are observed before and after treatment. Our formal derivations show that the linear probability DD will not only yield estimates that evolve with time since treatment but can also be opposite in sign from the true effect of treatment when the data are generated by a hazard process.

It is well-known that numerous issues arise when using logit or probit difference-in-differences procedures for a binary outcome (Ai and Norton 2003; Athey and Imbens 2006; Puhani 2012). Heckman (1996), among others, criticized these and other difference-in-differences procedures by noting the arbitrariness of such functional form assumptions. Nonlinearities are also key to our central result—that the linear probability DD will yield estimates that evolve with time since treatment if the data are generated by a hazard process. But as shown in the following, our formal result holds generally for single-decrement hazard processes that differ arbitrarily for treatment and controls.

To illustrate matters concretely, we reexamine findings by economists on unilateral and no-fault divorce (Friedberg 1998; Iyavarakul et al. 2011; Lee and Solon 2011; Wolfers 2006). In particular, we find that the increasing then decreasing pattern of effects noted by Wolfers (2006) can be generated by a time-invariant effect of treatment in a proportional hazard setting. More generally, our formal results emphasize that causal identification also requires assumptions, often implicit, on how the data are generated.

The article is organized as follows. We begin with the textbook case of two groups and two periods to derive the formal properties of the linear probability DD when the data are generated by a continuous-time hazard process. We then propose a three-step Cox estimation procedure for the proportional hazard DD that, to our knowledge, has not been previously discussed. Results from Monte Carlo simulations show that this estimator appears to perform well in practice. We then turn to a review of the empirical literature on unilateral and no-fault divorce, followed by results from a stylized example using empirical estimates of the baseline risk of divorce from the marital supplements to the June 1980, 1985, 1990, and 1995 Current Population Surveys. These results show that the increasing then decreasing pattern of estimates noted by Wolfers (2006) can be generated by a time-invariant effect of treatment in a proportional hazard setting. We conclude with some summary remarks.



**Fig. 1** The four components of the linear probability difference-in-differences estimator for two groups and two periods.

## Formal Derivations

For the textbook case of two groups and two periods, the linear probability DD can be written as

$$p_i = b_0 + b_1 \times g_i + b_2 \times I_i + dd_{lp} \times g_i \times I_i, \quad (1)$$

where  $p_i$  denotes the probability of the binary outcome,  $i$  indexes individuals, and  $g$  and  $I$  are dummy variables for the two groups and two periods, respectively, with  $g = 0$  referring to the control group,  $I = 1$  to the period in which treatment occurs, and  $dd_{lp}$  to what we will call the “linear probability DD.”

From Eq. (1), we have

$$p(g = 0, I = 0) = b_0$$

$$p(g = 0, I = 1) = b_0 + b_2$$

$$p(g = 1, I = 0) = b_0 + b_1$$

$$p(g = 1, I = 1) = b_0 + b_1 + b_2 + dd_{lp}$$

$$dd_{lp} = p(g = 1, I = 1) - b_0 - b_1 - b_2,$$

and hence that

$$dd_{lp} = [p(g = 1, I = 1) - p(g = 1, I = 0)] - [p(g = 0, I = 1) - p(g = 0, I = 0)]. \quad (2)$$

**Figure 1** depicts the four components of the linear probability DD.

Intuitively, the double difference in Eq. (2) can be seen as exploiting the (presumed) exogeneity of treatment and treatment timing while confronting the fact that neither controls nor treatments were randomly assigned. For concreteness, let the two groups be two U.S. states and the treatment be the introduction of unilateral and no-fault

divorce. Then to fix ideas, suppose that the true causal effect of unilateral and no-fault divorce is to increase divorce—that some in troubled marriages in the state that will receive treatment would not seek to divorce pretreatment but would do so posttreatment. But the fact that controls and treatments were not assigned at random means that it will not suffice to compare controls and treatments in the posttreatment period if, for example, the pretreatment level of divorce in the state that will later adopt unilateral and no-fault divorce is higher than in the state that will not. Thus, the double difference in Eq. (2) can be seen as accounting for this possibility in two (equivalent) ways. A first is to note that the DD in Eq. (2) adjusts for divorce trends by subtracting the difference in divorce between periods 2 and 1 for controls from the same quantity for treatments. A second is to note that Eq. (2) adjusts the naive comparison of controls and treatments in the posttreatment period by acknowledging that the nonrandom assignment of controls and treatments makes it likely that there were preexisting differences in divorce between controls and treatments in the pretreatment period.

The foregoing provides intuition into the logic of a difference-in-differences strategy, but Eq. (1) presumes that the linear probability DD is appropriate for a binary outcome such as divorce. But what if we were instead to view divorce as a continuous-time hazard process? We discuss the highly general case in which divorce in the two groups is given by two arbitrary hazard functions,  $r_{g=0}$  and  $r_{g=1}$ , but here we consider a proportional hazard DD that is the natural analog to the linear probability DD in Eq. (1):

$$r(t|t_0) = r_0(t - t_0) \exp[b_1 \times g_i + b_2 \times I_i(t) + dd_{hz} \times g_i \times I_i(t)]$$

or equivalently

$$\log r(t|t_0) = \log r_0(t - t_0) + b_1 \times g_i + b_2 \times I_i(t) + dd_{hz} \times g_i \times I_i(t), \quad (3)$$

where  $t$  denotes calendar time,  $t_0$  the calendar start of marriage,  $u = t - t_0$  marital duration,  $r_0$  the so-called baseline hazard,  $I(t)$  a time-varying dummy variable equal to 1 in the posttreatment period, and  $dd_{hz}$  the hazard difference-in-differences estimator.

If treatment and treatment timing are credibly exogenous for the linear probability DD, this too will hold equally for a hazard DD. Similarly, the same algebra relating Eqs. (1) and (2) can be used to reexpress  $dd_{hz}$  in Eq. (3) as a double difference, albeit for differences involving  $\log r$ :

$$dd_{hz} = [\log r(g = 1, I(t) = 1) - \log r(g = 1, I(t) = 0)] - [\log r(g = 0, I(t) = 1) - \log r(g = 0, I(t) = 0)]. \quad (4)$$

Could one modify the linear probability DD in Eq. (1) to mimic the proportional hazard DD in Eq. (3) by adding right-hand-side terms for marital duration? The answer is no, as can be seen by considering intervals of the form  $[u, u + \Delta]$ ,  $\Delta > 0$ . Then note that a key difference between the linear probability and hazard DD is that the latter compares the risk of divorce for controls and treatments in the interval  $[u, u + \Delta]$ , whereas the former does so for the probability of divorce. The issue then is that the linear probability DD ignores the fact that the comparison of divorce logically requires that divorce has not yet occurred as of the start of  $[u, u + \Delta]$ . By contrast, hazard analyses of divorce condition on those marriages that have survived as




Fig. 2 The probability of divorce pre- and posttreatment as a function of the survivor probability when data for divorce are generated by a continuous-time hazard process.

of the start of  $[u, u + \Delta]$ , with the classic life table taking  $\Delta$  to be some fixed positive constant and the continuous-time hazard taking the limit as  $\Delta \downarrow 0$ . Stated more formally, let  $u$  denote marital duration and  $U$  denote the random variable for duration; then the continuous-time hazard will be given by

$$r(u) = \lim_{\Delta \downarrow 0} \frac{\Pr(u < U \leq u + \Delta | U > u)}{\Delta} = \frac{f(u)}{S(u)},$$

where  $f(u)$  and  $S(u)$  denote the probability density and survivor probability functions for divorce. This textbook definition shows that  $r(u)$  differs from the unconditional probability of divorce by requiring that  $r(u)$  be defined only for marriages that survive to  $u$ .

The plausibility of causal claims from the linear probability DD requires the so-called “parallel trend” assumption—that net of level differences, controls and treatments are comparable to one another pretreatment and would continue to be comparable to one another were the treatment group not to have been treated in the post-treatment period. The corresponding comparability assumption for the proportional hazard DD is that controls and treatments share the common baseline hazard  $r_0(u)$ .

Finally, we note that a potential confound not well controlled by the linear probability DD is that treatment will occur at different marital durations for those from different marriage cohorts. It is thus natural in the two-group, two-period hazard case to adopt a cohort design in which the two groups are drawn from a single marriage cohort, with groups  $g = 0$  and  $g = 1$  thus beginning marriage at the same calendar time  $t_0$ .

We now turn to the central question posed in this study, which is what  $dd_{lp}$  estimates if divorces are in fact generated by a continuous-time hazard process. Let  $\tau$  denote the calendar time of treatment,  $t_0$  the calendar time at start of marriage, and  $[\tau_1, \tau]$  and  $[\tau, \tau_2]$  the pre- and posttreatment periods, respectively. Then as shown in Figure 2, the pre- and posttreatment probability of divorce, depicted by the two red

vertical bars, will be given by simple differences of the survivor probability, that is, the probability that the event of interest has not yet occurred.

Turning now to the general case in which the risk of divorce varies arbitrarily for treatments and controls, let  $r_{g=0}(u)$  and  $r_{g=1}(u)$  denote two arbitrary hazards functions of marital duration. Then  $S_g(u)$ , the probability of survival at duration  $u$  for group  $g$ , will be given by

$$S_g(t|t_0) = S_g(u) = \exp\left[-\int_0^{t-t_0} r_g(s)ds\right]. \quad (5)$$

Let  $dd_{lp(hz)}$  denote the linear probability difference-in-differences estimator if divorces are in fact generated by the arbitrary hazard functions  $r_{g=0}$  and  $r_{g=1}$ . Then, as shown in Figure 2,  $dd_{lp(hz)}$  will be given in expectation by

$$\begin{aligned} dd_{lp(hz)} &= \{E[p(g=1, I(t)=1)] - E[p(g=1, I(t)=0)]\} - \\ &\quad \{E[p(g=0, I(t)=1)] - E[p(g=0, I(t)=0)]\} \\ &= \{[S_{g=1}(\tau|t_0) - S_{g=1}(\tau_2|t_0)] - [S_{g=1}(\tau_1|t_0) - S_{g=1}(\tau|t_0)]\} - \\ &\quad \{[S_{g=0}(\tau|t_0) - S_{g=0}(\tau_2|t_0)] - [S_{g=0}(\tau_1|t_0) - S_{g=0}(\tau|t_0)]\}. \end{aligned} \quad (6)$$

Thus if divorces are generated by  $r_{g=0}$  and  $r_{g=1}$ , then Eq. (6) shows that the probability of divorce in the pre- and posttreatment periods will be a more complicated function than assumed by the linear probability DD. Recall that the probability of divorce for group  $g$  in period  $I$  under the linear probability DD in Eq. (1) is a simple function of the regression parameters  $b_0$ ,  $b_1$ , and  $dd_{lp}$ . By contrast, Eq. (6) shows that for the proportional hazard DD in Eq. (3), the probability of divorce for group  $g$  in period  $I$  will continue to be a function of  $b_0$ ,  $b_1$ , and  $dd_{hz}$ , but will also depend on: (1)  $t_0$ , the calendar time when marriage begins; (2)  $[\tau_1, \tau]$  and  $[\tau, \tau_2]$ , the intervals defining the pre- and posttreatment periods; and (3)  $S_g(t|t_0)$ , the survival probability for group  $g$ .<sup>1</sup>

*Theorem:* Let  $r_{g=0}(u)$  and  $r_{g=1}(u)$  be any two arbitrary hazard functions subject only to the condition that  $S_{g=0}$  and  $S_{g=1}$  be continuous and equal to 1 at the start of marriage; then  $dd_{lp(hz)}$  will evolve with time since treatment.

*Proof*

From Eq. (6), we have that  $dd_{lp(hz)}$  is a function of  $\tau_2$  via the two terms  $S_{g=0}(\tau_2|\tau)$  and  $S_{g=1}(\tau_2|\tau)$ . Then

$$dd_{lp(hz)} = S_{g=0}(\tau_2|t_0) - S_{g=1}(\tau_2|t_0) + c_1, \quad (7)$$

where

$$c_1 = [2S_{g=1}(\tau|t_0) - S_{g=1}(\tau_1|t_0)] - [2S_{g=0}(\tau|t_0) - S_{g=0}(\tau_1|t_0)] \quad (8)$$

<sup>1</sup> For an explicit expression for the proportional hazard  $dd_{lp(hz)}$  in Eq. (3), see the online supplement.

is a time-invariant constant for all  $\tau_2 > \tau$ . From Eq. (6), we have that  $dd_{lp(hz)}$  will evolve with time since treatment unless

$$S_{g=1}(\tau_2|t_0) - S_{g=0}(\tau_2|t_0) = c_2 \quad \forall \tau_2 > \tau \quad (9)$$

for some constant  $c_2$ . The condition in Eq. (9) requires that  $S_{g=0}$  and  $S_{g=1}$  be parallel for all  $\tau_2 > \tau$ , which will not hold in general except in three degenerate cases. To see this, first note that the condition that  $S_{g=0}$  and  $S_{g=1}$  will both equal 1 at  $t_0$  implies that  $S_{g=0}$  and  $S_{g=1}$  cannot be parallel for all  $t \geq t_0$  except in the degenerate case in which  $r_{g=0} = r_{g=1}$ .<sup>2</sup> However, Eq. (9) requires only that  $S_{g=0}$  and  $S_{g=1}$  be parallel for  $t > \tau$ . This too is highly restrictive, yielding two additional degenerative cases. We thus have that  $dd_{lp(hz)}$  will always evolve with time since treatment except in the following scenarios:

1. If  $r_{g=0}(t) = r_{g=1}(t) \quad \forall t \in [t_0, \infty]$  (no group differences and no effect of treatment);
2. If  $S_{g=0}(t) = S_{g=1}(t) = 0 \quad \forall t \in [\tau, \infty]$  (no posttreatment survivors); or
3. If  $r_{g=0}(t) = r_{g=1}(t) = 0 \quad \forall t \in [\tau, \infty]$  (no posttreatment events).

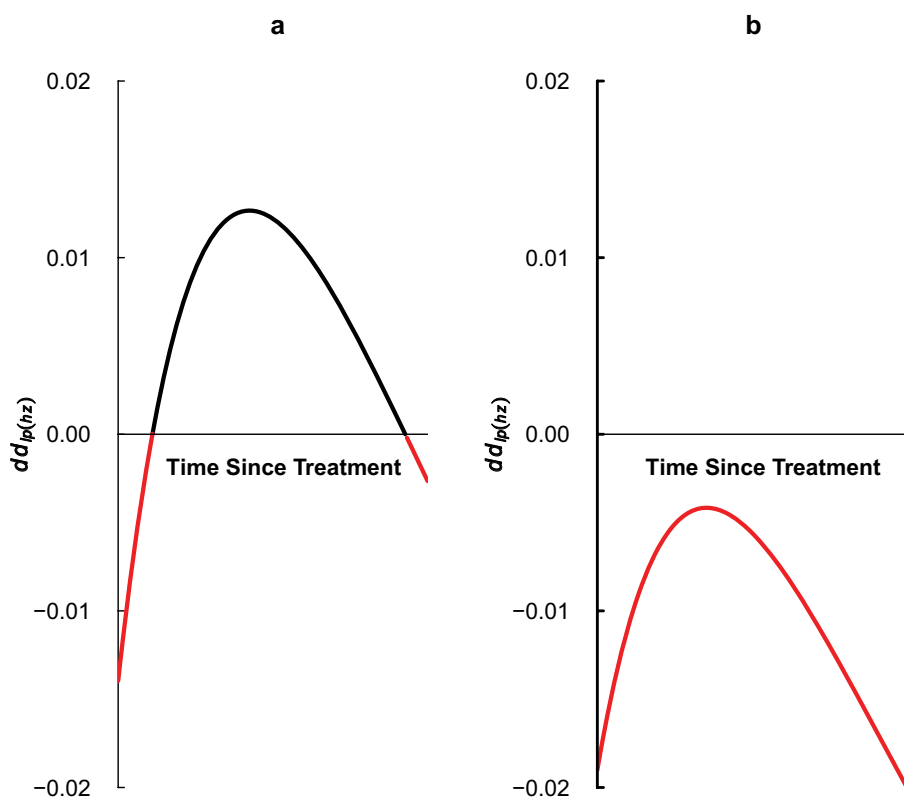
### Remarks

From Eq. (5), we have that the survivor probability will be a monotonically declining (more precisely, nonincreasing) function of marital duration; hence, the arithmetic difference of two such functions will also vary with marital duration, including marital durations in the posttreatment period. This implies that  $dd_{lp(hz)}$  is not a constant as assumed in the linear probability DD in Eq. (1), but will instead take values that evolve with time since treatment except in the foregoing three degenerate cases. The first is when divorce risks are identical for controls and treatments both pre- and posttreatment, which further implies no effect of treatment. The second involves so-called nondefective distributions for outcomes such as mortality in which all will experience the event of interest eventually. The third involves defective event distributions in which some will never experience the event of interest; examples include divorce, with some married couples never observed to divorce even when followed for a long time. The second and third degenerate cases then arise if the posttreatment period coincides with the period in which there are no survivors (Case 2) or no events (Case 3), respectively.

More fundamentally, Eq. (1) supposes that divorce is akin to a biased coin flip and hence that the effect of treatment is also akin to a biased coin flip. By contrast, Eq. (3) supposes that divorce is a continuous-time process involving the transition from an origin state (marriage) to a destination state (divorce). Thus under Eq. (3), divorces occur with exposure to risk, implying in turn that the probability of remaining married will be a nonincreasing function of marital duration.

*Corollary* (incorrect sign):  $dd_{hz} > 0 \not\Rightarrow dd_{lp(hz)} > 0$ .

<sup>2</sup> Requiring that  $S_{g=0}$  and  $S_{g=1}$  be equal to 1 at the start of marriage  $t_0$  rules out the case in which the hazard functions are identical for the two groups except at  $t_0$ , where the hazard for one group has a point mass spike such that  $S(t_0)$  is strictly less than 1 for one group but equal to 1 for the other group.



**Fig. 3** Behavior of  $dd_{ip(hz)}$  with time since treatment for an example in which the data are assumed to be generated according to the proportional hazard DD in Eq. (3). In both panels a and b, the baseline hazard  $r_0(u)$  is set equal to a constant  $\lambda = 0.008$ ; the regression coefficients  $b_1$ ,  $b_2$ , and  $dd_{hz}$  are assumed equal and set to 0.10; the calendar start of marriage,  $t_0$ , is set to 0; and the start of observation,  $\tau_1$ , is set to 24 months. In panel a, treatment begins at  $\tau = 60$  months of marital duration, while in panel b, treatment begins at  $\tau = 90$  months.

In Figure 3, we provide a simple example in which the linear probability and proportional hazard DD can be opposite in sign. In this example, we assume that divorces are generated by the proportional hazard specification in Eq. (3) and hence that  $dd_{hz}$  is the “true” causal effect of treatment. For computational convenience, we have taken the baseline hazard  $r_0(u)$  to be a constant  $\lambda$  equal to 0.008 divorces per month, thus yielding an exponential distribution for the timing of divorce, with the survivor function given, for example, by  $S(u) = \exp(-\lambda u)$  for  $g = 0$  and  $I(u) = 0$ . We also assume throughout that: (1) marriage begins at calendar time 0 ( $t_0 = 0$ ) for both controls and treatments; (2) observation begins two years after the start of marriage ( $\tau_1 = 24$  months); and (3) the regression coefficients in Eq. (3) take the values  $b_1 = b_2 = dd_{hz} = 0.10$ . Then given the foregoing, the only difference between panels a and b in Figure 3 is when treatment begins, at  $\tau = 60$  months and  $\tau = 90$  months, respectively, thus implying a shorter pretreatment period in panel a than in panel b.



As expected, panel a shows that  $dd_{lp(hz)}$  evolves with time since treatment, first rising to a peak and then declining, with  $dd_{lp(hz)}$  wrong in sign immediately after treatment and at longer durations since treatment, but correct in sign in between. Panel b provides a more extreme example, with  $dd_{lp(hz)}$  always negative and hence always wrong in sign.

In Figure 3, the values of  $dd_{lp(hz)}$  first rise then decline, but different values of  $b_1$ ,  $b_2$ ,  $dd_{hz}$ , or  $\lambda$  can imply patterns in which  $dd_{lp(hz)}$  appears to decrease then increase or increases (or decreases) monotonically with time since treatment. (Results are available upon request.)

## Estimation and Finite Sample Behavior of the Proportional Hazard DD

The formal derivations and examples discussed thus far provide a cautionary tale of what not to do—that is, how researchers can be badly misled by a linear probability DD when the data are in fact generated by a proportional hazard DD. Note also that the examples in Figure 3 hold in expectation and thus illustrate how  $dd_{lp(hz)}$  will evolve with time since treatment when the sample size of controls and treatments increases without bound. But what these derivations and stylized results do not speak to is how to estimate the hazard  $dd_{hz}$  in practice.

In this section, we propose a three-step estimator that shows how one can use the popular Cox model (Cox 1972) to obtain empirical estimates of  $dd_{hz}$ . We present Monte Carlo results that suggest that this procedure performs reasonably well in simulations in which the “true” value of  $dd_{hz}$  is known.

To fix ideas, we begin by first rearranging terms in Eq. (4) as


$$dd_{hz} = \left[ \log r(g=1, I(t)=1) - \log r(g=0, I(t)=1) \right] - \left[ \log r(g=1, I(t)=0) - \log r(g=0, I(t)=0) \right], \quad (10)$$

with the two bracketed terms referring to the posttreatment and pretreatment contrast between treatment and controls, respectively. This then suggests the following three-step estimation procedure:

1. Obtain an estimate of the first bracketed term using a Cox model and data from the posttreatment period.
2. Similarly, obtain an estimate of the second bracketed term using a Cox model and data from the pretreatment period.
3. Then estimate  $dd_{hz}$  via the arithmetic difference of these two estimates.

The rationale for using a Cox model in this procedure is that it yields asymptotically consistent and efficient estimates of proportional hazard regression parameters for an arbitrary baseline hazard  $r_0$ , thus allowing the analyst to obtain estimates in steps 1 and 2 without parametric assumptions about how divorce risks vary with marital duration in the pre- and posttreatment periods.

Deriving the asymptotic properties of the above three-step estimator is beyond the scope of this article, but we provide results from simulations that provide some sense of how well it may perform in practice. We begin with simulations that



**Fig. 4** Approximate normality by sample size of estimates of  $dd_{hz}$ . The simulated data were generated according to the proportional hazard DD in Eq. (3), with the baseline hazard  $r_0(u)$  set equal to a constant  $\lambda = 0.008$ ; the regression coefficients  $b_1$ ,  $b_2$ , and  $dd_{hz}$  are assumed equal and set to 0.10; the start of observation,  $\tau_1$ , is set to 24 months; and  $\tau$ , the start of treatment, is set to 60 months.

continue the example in panel a of Figure 3 in which the data are assumed to be generated by the proportional hazard DD in Eq. (3) with a constant baseline hazard  $\lambda = 0.008$  and with  $b_1 = b_2 = dd_{hz} = 0.10$ . We also assume a balanced design in which there are equal numbers in the control and treatment groups at the start of marriage. Figure 4 shows that estimates of  $dd_{hz}$  appear to follow a normal distribution that becomes more peaked as  $n$  increases, with the mean of the simulated estimates close to the true value  $dd_{hz} = 0.10$ .

Table 1 presents in tabular form the results from our Monte Carlo simulation for  $dd_{hz}$  corresponding to panels a and b in Figure 3. As in Figure 3, the two sets of estimates in Table 1 differ only in whether treatment begins at calendar time 60 or 90 months. The estimates of  $dd_{hz}$  exhibit some upward bias when treatment begins at 90 months, but the estimates in both cases remain within two standard deviations of 0.10, the true value of  $dd_{hz}$  as assumed in the Monte Carlo simulations.

Overall, Figure 4 and Table 1 provide suggestive evidence that a Cox model can be used to obtain reasonably precise estimates for the effect of treatment in a hazard difference-in-differences design in which  $n \geq 5,000$  in both the control and treatment groups.

**Table 1** Means and standard deviations of Cox estimates of  $dd_{hz}$  by sample size

Treatment at Calendar Time 60 Months			Treatment at Calendar Time 90 Months		
Mean	SD	$n$	Mean	SD	$n$
0.10053	0.07082	5,000	0.10520	0.05503	5,000
0.09944	0.04310	10,000	0.10486	0.04064	10,000
0.10061	0.03242	20,000	0.10471	0.02771	20,000

*Note:* The simulated data were generated according to the proportional hazard DD in Eq. (3), with the base-line hazard  $r_0(u)$  set equal to a constant  $\lambda = 0.008$ ; the regression coefficients  $b_1$ ,  $b_2$ , and  $dd_{hz}$  are assumed equal and set to 0.10; the start of observation,  $\tau_1$ , is set to 24 months; and  $\tau$ , the start of treatment, is set at 60 or 90 months.

**A Stylized Reexamination of Findings on Unilateral and No-Fault Divorce**

Although a large demographic literature has documented trends in, and the factors associated with, the dissolution of marital unions, it has been economists who have posed the causal question of whether the introduction of unilateral and no-fault divorce laws caused an increase in divorce (Friedberg 1998; Iyavarakul et al. 2011; Lee and Solon 2011; Wolfers 2006), with two of these studies (Friedberg 1998; Wolfers 2006) appearing in the *American Economic Review*, the flagship journal of the American Economic Association.

Because laws at the state level govern divorce in the United States, these studies all relied on state-level data and a linear probability difference-in-differences design to identify the causal effect of the shift from laws requiring the mutual consent of both spouses to a legal standard allowing one spouse to seek divorce on grounds such as the irretrievable breakdown of the marriage, spousal incompatibility, or irreconcilable differences. These studies also considered the possibility that trends in divorce may have varied considerably across states during the period when the shift to unilateral and no-fault divorce was taking place, thus complementing empirical studies documenting historical trends in divorce in the United States as a whole (Cherlin 1991; Preston and McDonald 1979).

Friedberg (1998) was the first to use a linear probability DD to estimate the effect of unilateral and no-fault divorce. Friedberg used divorce registers for the period 1968–1988 to construct pre- and posttreatment panel data on the annual number of divorces per 1,000 persons for the 50 states and the District of Columbia. In models specifying state and calendar year fixed effects and state-specific linear and quadratic trends, Friedberg’s DD estimates implied increases of between 0.441 and 0.447 divorces per 1,000, or a roughly 9.5% and 9.7% increase, respectively, on a baseline of 4.6 divorces per 1,000. Friedberg obtained similar estimates when distinguishing between the stringency of unilateral and no-fault divorce decrees, obtaining estimates implying increases of 9.7% to 11.9%. These results led Friedberg to conclude that “the effect of unilateral divorce on divorce behavior was permanent, not temporary” (Friedberg 1998:608).

Wolfers (2006) provided both a replication and critique of Friedberg. Analyzing data generously provided to him by Friedberg, he replicated her estimate of a roughly

10% increase in divorce when using Friedberg's preferred linear probability specification. He then raised the possibility that married couples might respond dynamically to the introduction of unilateral and no-fault divorce. To investigate this empirically, he modified the linear probability DD specified by Friedberg to allow the effect of treatment to vary with time since treatment. In analyses of data in which he extended Friedberg's panel to cover the period 1958–1967, he obtained estimates that first increased then decreased with time since treatment. These results led Wolfers to conclude, in contrast to Friedberg, that there “is no evidence that this rise in divorce is persistent” but also that his results “suggest—somewhat puzzlingly—that 15 years after reform the divorce rate is lower as a result of the adoption of unilateral divorce, although it is hard to draw any strong conclusions about long-run effects” (Wolfers 2006:1802).<sup>3</sup>

In a study that anticipated in part some of the issues we raise, Iyavarakul et al. (2011) proposed a theoretical model to account for the apparent variation in the estimated effect of unilateral divorce with time since treatment. In their model, the time-varying effect of no-fault divorce is due to the forward-looking behavior of three distinct groups of married couples: (1) those who marry and divorce prior to treatment; (2) those who marry after treatment and whose selection into marriage was therefore influenced by treatment; and (3) those who marry before, but remain married after, treatment and who are therefore “surprised” by treatment. Their model thus implies that selection into marriage will differ for these three groups and that the effect of treatment will likewise differ across groups. A core element of their behavioral model thus concerns the behavior of successive marriage cohorts; however, their analyses rely on the same aggregate panel data assembled by Friedberg and Wolfers, thus leading them to model the outcome as the probability of divorce in a given state, year, and treatment by group cell.

Lee and Solon (2011) reanalyzed the data used by Wolfers and concluded that the increasing then decreasing pattern of estimates reported by Wolfers is highly sensitive to functional form, autocorrelation, weighting, and other issues. For example, they found little effect of unilateral divorce when analyzing the natural logarithm of divorces per capita as well as substantial first- and higher order autocorrelations among the residuals in the weighted least-squares specifications used by Wolfers. They concluded that “the true impact of unilateral divorce laws remains unclear” (Lee and Solon 2011).

The data analyzed in these studies were either Friedberg's original data or Wolfers' additions to the Friedberg data; hence, the number of divorces per 1,000 persons in a given state and calendar year is the analytic outcome used in all four studies.<sup>4</sup>

<sup>3</sup> To motivate dynamic response to treatment, Wolfers (2006:1806) sketched a simple model that posited heterogeneity in the “compatibility” of married couples and in which “under consent divorce laws, the 20 percent most incompatible matches dissolve [while] under unilateral divorce, this rises to 20.4 percent.” He also provided an important and insightful discussion of issues from a stock and flow point of view, although data limitations precluded him from conducting such a stock and flow analysis. Formulating divorce as a hazard process provides a natural framework for modeling outflows from the stock of marriage due to divorce. See, for example, the hazard analyses in Preston and McDonald (1979) for outflows from marriage due to death and divorce, and Klerman and Haider (2004) for outflows from the stock of welfare recipients due to policies placing time limits on the receipt of welfare.

<sup>4</sup> Lee and Solon (2011) analyze both divorces per capita and log divorces per capita.

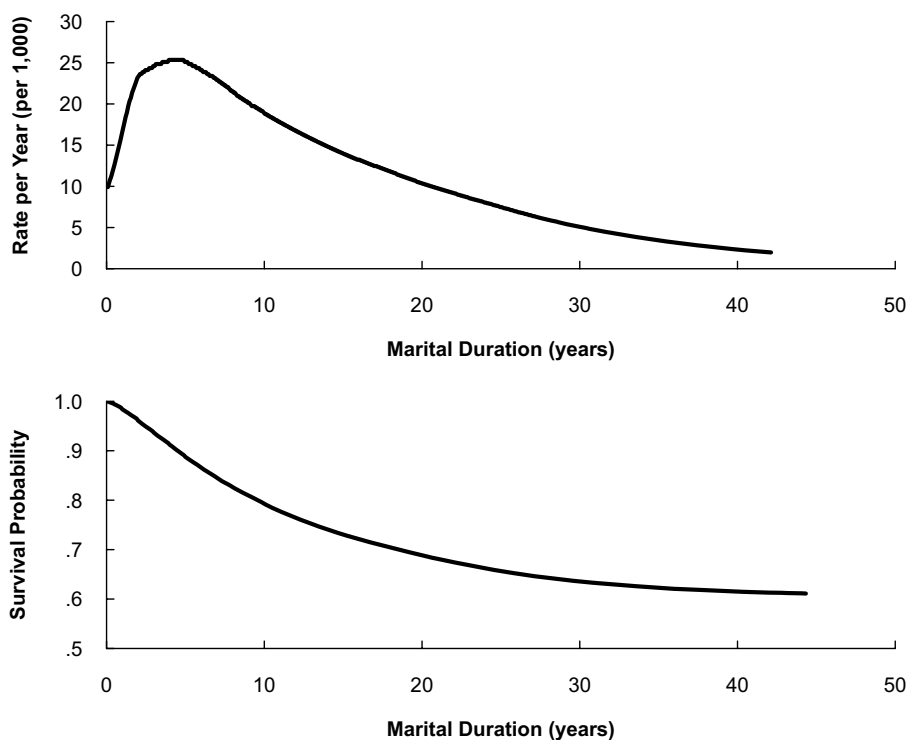
As acknowledged by Friedberg (1998:611), this measure differs from a period rate such as the number of divorces per 1,000 marriages. Note that divorces per 1,000 persons will be downwardly biased relative to divorces per 1,000 marriages because the denominator for the former will be substantially larger than the denominator for the latter.

A more fundamental issue is the precise question being posed when analyzing divorce probabilities using measures such as divorces per capita versus the risk of divorce using measures such as divorces per 1,000 marriages. In the foregoing DD analyses, the number of divorces per capita in state  $s$  and calendar year  $t$  would appear to be interpretable in much the same way as a period divorce rate given by divorces per 1,000 marriages, with both ostensibly capturing change in divorce behavior. But because the two measures differ by the multiplicative factor of marriages per capita, trends in divorces per capita can occur even if there is no trend in actual divorce behaviors if there are trends in either the numerator or denominator of marriages per capita. These issues do not arise in hazard analyses that condition on marriage, hence restricting attention to those who are actually at risk of divorce.

We now turn to a stylized reexamination of findings on unilateral and no-fault divorce. Estimation of the proportional hazard DD places far greater demands on data than the linear probability DD, with the proportional hazard DD requiring not only marital histories but also residential histories allowing one to track changes in the state in which a married couple resides. The Panel Study of Income Dynamics is one possibility, but Friedberg (1998:610) noted that the number of married couples is not large enough to provide sufficient statistical power to obtain reasonably precise estimates, with Friedberg's assertion also consistent with the Monte Carlo simulation results reported in Table 1. We instead analyze data from the marital supplement to the June 1980, 1985, 1990, and 1995 Current Population Surveys (CPS), which provide large samples when pooled.

Our use of the term "stylized" is intended to flag the fact that the CPS lacks a residential history, providing only the state of residence at the time of CPS survey. These data limitations thus prevent us from a full analysis contrasting estimates from the linear probability and proportional hazard DD. In the analyses that follow, we instead present stylized results for the textbook two-group, two-period case. These analyses combine empirical estimates of the baseline risk of divorce obtained from the June CPS with posited values of the regression coefficients in the proportional hazard DD, with the posited values chosen to be consistent with those reported by Friedberg and Wolfers.

The retrospective marital histories in the CPS were obtained from married females aged 15 or older and never-married females aged 18 or older. Respondents were asked about the number of marriages, which was then followed by data on the first two and most recent marriage. The resulting marital histories thus provide the calendar year and month when a marriage began and, if a marriage ended, the dates of widowhood, separation, or divorce. In the June 1995 supplement, marital histories were obtained for the first three and most recent marriages. The pooled CPS sample contains 201,033 female respondents, which we then restricted by dropping never-married females ( $n=45,881$ ) and a small number of cases with missing data ( $n=2,477$ ), yielding an analytic sample of 152,675 ever-married females. We then used these data to construct marital histories providing data on the duration of




**Fig. 5** Divorce risk (upper panel) and survivor probability (lower panel) by marital duration among U.S. women. *Source:* June 1980, 1985, 1990, and 1995 Current Population Survey.

marriage to the nearest month at divorce. We censored marriages at CPS survey, widowhood, or separation if the respondent reported a separation but no subsequent divorce. The resulting data contain 185,047 marriages and 47,655 divorces (183,781.6 and 47,760.6, respectively, when weighted). The data cover the period both before and after the adoption of unilateral and no-fault divorce by states, with marriage cohorts that began marriages as early as 1928 and as late as 1995.

Figure 5 reports nonparametric estimates showing how divorce risks (upper panel) and survivor probabilities (lower panel) vary with marital duration. Estimates of divorce risks were obtained using a procedure described in Wu (1989); estimates of survivor probabilities were obtained using the Kaplan–Meier estimator (Kaplan and Meier 1958). Divorce risks increase rapidly at early marital durations then decline at later marital durations. Divorce risks peak at around 4.5 years of marriage at a level of roughly 25 divorces per 1,000 marriages per year. Survivor probabilities decline monotonically with marital duration, with roughly four in 10 divorces occurring in these marriage cohorts.

We now turn to our stylized two-group, two-period example, in which we suppose that a researcher reports findings from the linear probability DD when divorce is in fact generated by the proportional hazard process in Eq. (3). As noted, the resulting analysis is stylized, in which we (1) use the CPS data to obtain empirical estimates of  $r_0(u)$ , the baseline risk of divorce, but (2) posit hypothesized values for the regression coefficients  $b_1$ ,  $b_2$ , and  $dd_{hz}$  in the proportional hazard DD in Eq. (3).



**Fig. 6** Behavior of  $dd_{lp(hz)}$  with time since treatment. The stylized example assumes  $b_1 = b_2 = 0.10$  in Eq. (3) and pre- and posttreatment observation periods equal to  $[60, 120]$  and  $[120, 480]$ , respectively. Empirical estimates of baseline divorce risks were obtained from retrospective marital histories reported by female respondents in the June 1980, 1985, 1990, and 1995 Current Population Survey.

To estimate  $r_0(u)$ , we used a highly flexible piecewise splined Gompertz specification with knots at 18, 24, 48, 72, 96, 120, and 180 months of marital duration, thus yielding a piecewise linear spline for  $\log r_0(u)$ . We then posited the following: (1) that  $b_1 = 0.1$ , consistent with group differences in which higher levels of divorce were observed in states that initially adopted unilateral and no-fault divorce; (2) that  $b_2 = 0.1$ , consistent with the increasing trend in divorce during this period; (3) that pretreatment observation begins at 5 years (60 months) of marital duration; and (4) that treatment occurs at 10 years (120 months) of marital duration.

Figure 6 depicts how  $dd_{lp(hz)}$  evolves with time since treatment when the data are generated by the proportional hazard process in Eq. (3). The rising then declining values of  $dd_{lp(hz)}$  require only that  $b_1 > 0$  and  $b_2 > 0$  and are otherwise robust to the values specified in (1)–(4). (These results are available upon request.) As expected, time-invariant values of  $dd_{hz}$  imply estimates of  $dd_{lp(hz)}$  that vary with time since treatment, first rising then declining with time since treatment. For  $dd_{hz} = 0.0$ , we see that  $dd_{lp(hz)}$  takes values that are small in magnitude, with  $dd_{lp(hz)}$  negative initially, then positive,



and then negative with time since treatment. This inverted U-shape pattern becomes more pronounced as  $dd_{hz}$  takes larger positive values, thus potentially tempting those employing Eq. (1) to interpret the resulting patterns as increasingly credible evidence of dynamic response to treatment such as “pent-up” demand (Wolfers 2006:1806). These results thus suggest that for plausible values of group difference ( $b_1 = 0.1$ ) and trends in divorce ( $b_2 = 0.1$ ), time-invariant values of the proportional hazard DD coefficient  $dd_{hz}$  will generate values of  $dd_{lp(hz)}$  that yield the qualitative pattern of results reported by Wolfers (2006).

Note that the curves for  $dd_{lp(hz)}$  take the same negative value at  $\tau$ , the time at start of treatment, even when  $dd_{hz}$  takes different values. This follows from  $b_1 > 0$ , which dictates that the pretreatment probability of divorce is higher for  $g = 1$  than for  $g = 0$ , thus yielding a common negative intercept  $p(g = 0, t = \tau) - p(g = 1, t = \tau)$ . Larger positive values of  $b_1$  or wider pretreatment intervals will increase the absolute magnitude of this negative intercept.

A standard “frailty” hypothesis is that some married couples will be more divorce-prone than others, a possibility also suggested by the nonparametric survivor probabilities in Figure 5. Heterogeneity in “divorce proneness” would, in turn, imply a changing composition in a marriage cohort as divorces occur to the more divorce-prone, leaving a surviving stock of marriages that will be increasingly less divorce-prone. For the linear probability DD, this poses a threat to the credibility of causal claims because compositional change will act as an unobserved time-varying confound. Note, however, that if the goal is to obtain a credibly causal estimate of treatment on the treated, a potential confound that varies with marital duration can be treated in a hazard setting as a nuisance function in the sense of Cox (1972).

## Discussion

For those wishing to use a difference-in-differences design to analyze a binary outcome, we argue that what this study implies is straightforward. If the binary outcome is best viewed as something akin to a biased coin flip, then a linear probability DD may well be appropriate. But if the binary outcome is best viewed as a single-decrement, continuous-time process involving the transition from one discrete state to another, then the linear probability DD should be avoided and a hazard DD used instead. For some, this conclusion and our formal results may be seen as the unsurprising consequence of model misspecification. Still, that a binary outcome generated by a hazard process differs fundamentally from a biased coin flip—something long understood in the field of demography—is perhaps less well recognized, at least by some in other disciplines. But perhaps most importantly, our results emphasize that for a binary outcome, yet another necessary aspect of causal identification are assumptions, often implicit, about how the data are generated.

We have restricted attention to the textbook case of two groups and two periods, but different states adopted unilateral and no-fault divorce in different years, thus requiring DD procedures that generalize to multiple groups and multiple pre- and posttreatment periods. For such real-world data, standard practice has been to specify state and calendar year fixed effects, and the analogous proportional hazard DD would be to add a third set of fixed effects for marriage cohort. However, important



recent developments show that the resulting DD regression coefficient will not in general give the desired causal estimate for the average effect of treatment on the treated, which instead can be shown to be equal to a weighted average of all possible two-group, two-period DDs (Callaway and Sant'Anna 2021; de Chaisemartin and D'Haultfœuille 2022; Goodman-Bacon 2021). Although these results can be seen as redirecting attention back to the core role played by the textbook two-group, two-period DD considered in this study, an unanswered question for future research is whether a similar result holds for the proportional hazard DD when there are multiple groups and periods.

The popularity of DD procedures stems in no small part from the wide range of data that can be used, including not only panel data following individuals or units over time but also repeated cross sections for outcomes observed at the aggregate level for geographic units such as states or counties. Thus in the case of unilateral and no-fault divorce, economic studies to date have used a linear probability DD to analyze panel data on divorces measured at the state level. By contrast, the hazard DD makes far greater data demands, requiring individual-level data containing marriage histories on when a marriage began and the date of divorce if divorce occurred, but also a residential history in the state in which a married couple resided during the course of their marriage.

To date, there have been notable points of disagreement in economic studies of unilateral and no-fault divorce, including (1) whether effects are positive and persistent (Friedberg 1998), (2) whether effects are positive but subject to dynamic response (Iyavarakul et al. 2011; Wolfers 2006), or (3) whether estimates are too fragile to warrant any firm conclusion (Lee and Solon 2011). But these studies share a common albeit implicit assumption—that biased coin flips generate both the outcome and effect of treatment. We contribute to these debates by showing that the rising then declining effect of unilateral and no-fault divorce noted by Wolfers (2006) can be generated by a time-invariant difference-in-differences coefficient in a proportional hazard setting. Thus, like Lee and Solon (2011), we conclude that the true impact of unilateral and no-fault divorce laws remains unclear, but we reach this conclusion on fundamentally different grounds. ■

**Acknowledgments** Early versions of this article were presented at the 2019 Summer Research Workshop, Institute for Research on Poverty, University of Wisconsin–Madison, and as the Duncan Lecture at the 2019 annual meeting of the American Sociological Association. We thank the *Demography* reviewers, Richard Breen, Siwei Cheng, Rajeev Dehejia, Andrew Goodman-Bacon, William Greene, Michael Hout, Nicholas Mark, Robert Moffitt, Jeff Smith, Ross Stolzenberg, and Chris Taber for comments on early drafts.

## References

- Ai, C., & Norton, E. C. (2003). Interaction terms in logit and probit models. *Economics Letters*, 80, 125–129.
- Ashenfelter, O., & Card, D. (1985). Using the longitudinal structure of earnings to estimate the effect of training programs. *Review of Economics and Statistics*, 67, 648–660.
- Athey, S., & Imbens, G. W. (2006). Identification and inference in nonlinear difference-in-differences models. *Econometrica*, 74, 431–497.

- Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225, 200–230.
- Card, D., & Krueger, A. B. (1994). Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania. *American Economic Review*, 84, 772–793.
- Cherlin, A. J. (1991). *Marriage, divorce, remarriage*. Cambridge, MA: Harvard University Press.
- Cox, D. R. (1972). Regression models and life-tables. *Journal of the Royal Statistical Society: Series B*, 34, 187–220.
- de Chaisemartin, C., & D'Haultfoeuille, X. (2022). *Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: A survey* (NBER Working Paper 29691). Cambridge, MA: National Bureau of Economic Research. <https://doi.org/10.3386/w29691>
- Friedberg, L. (1998). Did unilateral divorce raise divorce rates? Evidence from panel data. *American Economic Review*, 88, 608–627.
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225, 254–277.
- Heckman, J. J. (1996). Comment on “Labor supply and the Economic Recovery Tax Act of 1981,” by Nada Eissa. In M. Feldstein and J. M. Poterba (Eds.), *Empirical foundations of household taxation*. Chicago: University of Chicago Press.
- Iyavarakul, T., McElroy, M. B., & Staub, K. (2011). *Dynamic optimization in models for state panel data: Model of divorce laws on divorce rates* (ERID Working Paper No. 140). Durham, NC: Duke University, Department of Economics. Retrieved from <http://ssrn.com/abstract=2210408>
- Kaplan, E. L., & Meier, P. (1958). Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association*, 53, 437–481.
- Klerman, J. A., & Haider, S. J. (2004). A stock-flow analysis of the welfare caseload. *Journal of Human Resources*, 39, 865–886.
- Lee, J. Y., & Solon, G. (2011). The fragility of estimated effects of unilateral divorce laws on divorce rates. *B.E. Journal of Economic Analysis & Policy*, 11, 49. <https://doi.org/10.2202/1935-1682.2994>
- Preston, S. H., & McDonald, J. (1979). The incidence of divorce within cohorts of American marriages contracted since the Civil War. *Demography*, 16, 1–25.
- Puhani, P. A. (2012). The treatment effect, the cross difference, and the interaction term in nonlinear “difference-in-differences” models. *Economics Letters*, 115, 85–87.
- Wolfers, J. (2006). Did unilateral divorce laws raise divorce rates? A reconciliation and new results. *American Economic Review*, 96, 1802–1820.
- Wu, L. L. (1989). Issues in smoothing empirical hazards. In C. C. Clogg (Ed.), *Sociological methodology, 1989* (Vol. 19, pp. 127–159). Washington, DC: American Sociological Association.

---

Lawrence L. Wu (corresponding author)  
[lawrence.wu@nyu.edu](mailto:lawrence.wu@nyu.edu)

Wu • Department of Sociology, New York University, New York, NY, USA; <https://orcid.org/0000-0003-2301-5136>

Wen • Department of Political and Social Change, Australian National University, Canberra, Australia; <https://orcid.org/0000-0001-5075-2831>