

**HAZARD VS. LINEAR PROBABILITY
DIFFERENCE-IN-DIFFERENCES
ESTIMATORS FOR DEMOGRAPHIC PROCESSES**

Lawrence L. Wu
New York University

Fangqi Wen
University of Oxford

April 2021

Earlier versions of this paper 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 Meetings of the American Sociological Association, New York, NY. We thank 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 helpful comments and discussions. Direct all correspondence to Lawrence L. Wu, Department of Sociology, 295 Lafayette Street, Puck Building, 4th floor, New York University, New York, NY 10012-9605, lawrence.wu@nyu.edu.

**HAZARD VS. LINEAR PROBABILITY
DIFFERENCE-IN-DIFFERENCES
ESTIMATORS FOR DEMOGRAPHIC PROCESSES**

ABSTRACT

(132 words)

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 derive formal results for the textbook case of two groups and two periods in which the control and treatment groups are observed pre- and post-treatment. We 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, marriage cohort, unilateral and no-fault divorce

1. 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 geographical 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 show that for the textbook case of two groups and two periods, 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.¹ To illustrate matters concretely, we reexamine findings by economists on unilateral and no-fault divorce (Friedberg 1998; Wolfers 2006; Iyavarakul, McElroy, and Staub 2011; Lee and Solon 2011). In particular, we find that the increasing then decreasing pattern of estimates noted by Wolfers (2006)

¹It 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 2003; Puhani 2012). Heckman (2006), among others, has 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 in fact generated by a hazard process. But this result does not rely on specific functional forms, holding instead for hazards that can differ arbitrarily for treatment and controls.

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 paper 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 in fact generated by a continuous-time hazard process. We then review 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.

2. 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 where 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 term the “linear probability DD”.

From (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

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

See Figure 1, which depicts the four components of the linear probability DD.

[Figure 1 about here]

Intuitively, the double difference in (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 pre-treatment but would do so post-treatment. 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 post-treatment period if, for example, the pre-treatment 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 (2) can be seen as accounting for this possibility in two (equivalent) ways. A first is to note that the DD in (2) adjusts for divorce trends by subtracting the pre- vs. post increase in divorce for controls from the pre- vs. post- increase in divorce for treatments. A second is to note that (2) adjusts the naive comparison of controls and treatments in the post-treatment period by acknowledging that the non-random assignment of controls and treatments makes it likely that there were pre-existing differences in divorce between controls and treatments in the pre-treatment period.

The above provides intuition into the logic of a difference-in-differences strategy, but (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 below in which divorce in the two groups are 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 equation (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 post-treatment period, and dd_{hz} the hazard difference-in-differences estimator. Estimation of equation (3) is straightforward with, for example, the Cox model (Cox 1972) providing consistent and asymptotic efficient estimates of b_1 , b_2 , and dd_{hz} for arbitrary r_0 .

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 equations (1) and (2) can be used to re-express dd_{hz} in (3) as a double difference, albeit for differences involving $\log r$:

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

Could one modify the linear probability DD in (1) to mimic the proportional hazard DD in (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 of the start of $[u, u + \Delta]$, with the classic life table taking Δ 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 at divorce; then the continuous-time hazard will be given by

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

where $f(u)$ and $S(u)$ denote the probability density and survivor probability functions for divorce. The above 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 pre-treatment 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 τ 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 post-treatment periods, respectively. Then as shown in Figure 2, the pre- and post-treatment probability of divorce, depicted by the two 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.

[Figure 2 about here]

Turning now to highly 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:

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

$$\{[S_{g=0}(\tau|t_0) - S_{g=0}(\tau_2|t_0)] - [S_{g=0}(\tau_1|t_0) - S_{g=0}(\tau|t_0)]\}. \quad (6)$$

Thus if divorces are generated by $r_{g=0}$ and $r_{g=1}$, then equation (6) shows that the probability of divorce in the pre- and post-treatment 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 (1) is a simple function of the regression parameters b_0 , b_1 and dd_{lp} . By contrast, equation (6) shows that for the proportional hazard DD in (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: (i) t_0 , the calendar time when marriage begins; (ii) $[\tau_1, \tau]$ and $[\tau, \tau_2]$, the intervals defining the pre- and post-treatment periods; and (iii) $S_g(t|t_0)$, the survival probability for group g .²

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 equation (6), we have that $dd_{lp(hz)}$ is a function of τ_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

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

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

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

for some constant c_2 . The condition in equation (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

²For an explicit expression for the proportional hazard $dd_{lp(hz)}$ in (3), see Appendix.

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

Remarks: From equation (5), we have that the survivor probability will be a monotonically declining (more precisely, non-increasing) function of marital duration; hence, the arithmetic difference of two such functions will also vary with marital duration, including marital durations in the post-treatment period. This implies that $dd_{lp(hz)}$ is not a constant as assumed in the linear probability DD in (1), but will instead take values that evolve with time since treatment except in the above three degenerate cases. The first is when divorce risks are identical for controls and treatments both pre- and post-treatment, which further implies no effect of treatment. The second involves so-called non-defective 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 post-treatment period coincides with the period in which there are no survivors (Case 2) or no events (Case 3), respectively.

More fundamentally, equation (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, equation (3) supposes that divorce is a continuous-time process involving the transition from an origin state (marriage) to a destination state (divorce). Thus under equation (3), divorces occur with exposure to risk, implying in turn that the probability of remaining married will be a non-increasing function of marital duration.

Corollary 1 (incorrect sign): $dd_{hz} > 0 \not\Rightarrow dd_{lp(hz)} > 0$

³Requiring 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.

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 (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 λ 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: (i) that marriage begins at calendar time 0 ($t_0 = 0$) for both controls and treatments; (ii) that observation begins two years after the start of marriage ($\tau_1 = 24$ months); and (iii) that the regression coefficients in (3) take the values $b_1 = b_2 = dd_{hz} = 0.10$. Then given the above, the only difference between Panels A and B is when treatment begins, at $\tau = 60$ months in Panel A vs. $\tau = 90$ months in Panel B, thus imply a shorter pre-treatment period in Panel A than in Panel B.

[Figure 3 about here]

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 λ can imply patterns in which $dd_{lp(hz)}$ appears to decrease then increase or increases (or decreases) monotonically with time since treatment. (Results available upon request.)

Finally, note that the $dd_{lp(hz)}$ in Figure 3 follow from equation (6) and thus hold in expectation. Put another way, if the binary outcome of interest is generated by the hazard process in (3), the substantive implication of Panel B is that the researcher employing the linear probability DD in (1) can obtain the wrong sign for the causal effect of treatment even when the sample analyzed is arbitrarily large. We thus conclude that for a binary outcome, often implicit assumptions about how the data are generated are an important and necessary component of causal identification.

3. 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 examined the causal question of whether the

introduction of unilateral and no-fault divorce laws led to an increase in divorce (Friedberg 1998, Wolfers 2006, Iyavarakul, McElroy and Staub 2011, and Lee and Solon 2011), 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 (Preston and McDonald 1979; Cherlin 1991).

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 post-treatment 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 percent increase 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 percent. These results led Friedberg to conclude that “the effect of unilateral divorce on divorce behavior was permanent, not temporary” (p. 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” (p. 1802).⁴

In a paper that anticipates in part some of the issues we raise, Iyavarakul, McElroy and Staub (2011) propose 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 find 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 conclude that “the true impact of unilateral divorce laws remains unclear”.

The data analyzed in the above studies were either Friedberg’s original data or Wolfer’s 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.⁵ As acknowledged by Friedberg (footnote 8, p. 611), this measure differs from a period rate such as the number of divorces per 1,000 marriages. Note that divorces per 1,000

⁴To motivate dynamic response to treatment, Wolfers (2006) 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” (p. 1806). 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, e.g., 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.

⁵Lee and Solon analyze both divorces per capita and log divorces per capita.

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 vs. the risk of divorce using measures such as divorces per 1,000 marriages. In the above 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 (PSID) is one possibility, but Friedberg (1998, p. 610, footnote 6) notes that the relatively small samples in the PSID do not provide sufficient statistical power to obtain reasonably precise estimates. 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, and/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 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 contains 185,047 marriages and 47,655 divorces (183,781.6 and 47,760.6 weighted).

Figure 4 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 1957). 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 4 in 10 divorces occurring in these marriage cohorts.

[Figure 4 about here]

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 equation (3). As noted above, the resulting analysis is stylized, in which we: (i) use the CPS data to obtain empirical estimates of $r_0(u)$, the baseline risk of divorce, but (ii) posit hypothesized values for the regression coefficients b_1 , b_2 , and dd_{hz} in the proportional hazard DD in equation (3).

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: (i) that $b_1 = 0.1$, consistent with group differences in which higher levels of divorce were observed in states that initially adopted universal and no-fault divorce; (ii) that $b_2 = 0.1$, consistent with the increasing trend in divorce during this period; (iii) that pre-treatment observation begins at 5 years (60 months) of marital duration; and (iv) that treatment occurs at 10 years (120 months) of marital duration.

Figure 5 depicts how $dd_{lp(hz)}$ evolves with time since treatment when the data are generated by the

proportional hazard process in equation (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 (i)–(iv) (results 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 equation (1) to interpret the resulting patterns as increasingly credible evidence of dynamic response to treatment such as “pent-up” demand (Wolfers 2006, p. 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).

[Figure 5 about here]

Note that the curves for $dd_{lp(hz)}$ take the same negative value at τ , the time at start of treatment even when dd_{hz} takes different values. This follows from $b_1 > 0$, which dictates that the pre-treatment 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 pre-treatment 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 4. 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).

4. Discussion

For those wishing to use a difference-in-differences design to analyze a binary outcome, what this study implies is, we argue, 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 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 post-treatment 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 does not in fact 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 (Goodman-Bacon 2019; de Chaisemartin and DHaultfuille 2019). 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 geographical 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 on 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 (Wolfers 2006; Iyavarakul, McElroy, and Staub 2011), 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-difference 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 than Lee and Solon.

REFERENCES

- Ai, Chunrong, and Edward C. Norton. 2003. "Interaction Terms in Logit and Probit Models." *Economic Letters* 80: 125–129.
- Ashenfelter, Orley; and David Card. 1985. "Using the Longitudinal Structure of Earnings to Estimate the Effect of Training Programs." *The Review of Economics and Statistics* 67(4): 648–660.
- Athey, Susan, and Guido W. Imbens. 2006. "Identification and Inference in Nonlinear Difference-in-Differences Models." *Econometrica* 74(2): 431–497.
- Card, David, and Alan B. Krueger. 1994. "Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania." *The American Economic Review* 84(4): 772–793.
- Cherlin, Andrew J. 1981. *Marriage, Divorce, Remarriage*. Cambridge, MA: Harvard University Press.
- Cox, D. R. 1972. "Regression Models and Life Tables (with discussion)." *Journal of the Royal Statistical Society* 4B(2): 187–220.
- de Chaisemartin, Clement, and Xavier D'Haultfoeuille. 2019. "Two-Way Fixed Effects Estimators With Heterogeneous Treatment Effects." Working Paper.
- Friedberg, Leora. 1998. "Did Unilateral Divorce Raise Divorce Rates? Evidence from Panel Data." *American Economic Review*. 88(3): 608–27.
- Goodman-Bacon, Andrew. 2019. "Difference-In-Differences With Variation in Treatment Timing." Working Paper.
- Iyavarakul, Tongyai, Marjorie B. McElroy, and Kalina Staub. 2011. "Dynamic Optimization in Models for State Panel Data: Model of Divorce Laws on Divorce Rates." ERID Working Paper No. 140, Duke University, <http://ssrn.com/abstract=2210408>
- Kaplan, E. L., and Paul Meier. 1958. "Nonparametric Estimation from Incomplete Observations." *Journal of the American Statistical Association* 53(282): 437–81.
- Klerman, Jacob Alex, and Steven J. Haider. 2004. "A Stock-Flow Analysis of the Welfare Caseload." *Journal of Human Resources* 39(4): 865–86.

- Lee, Jin Young, and Gary Solon. 2011. "The Fragility of Estimated Effects of Unilateral Divorce Laws on Divorce Rates." *B.E. Journal of Economic Analysis and Policy* (Contributions) 11(1): Article 49.
- Preston, Samuel H., and John McDonald. 1979. "The Incidence of Divorce within Cohorts of American Marriages Contracted since the Civil War." *Demography* 16(1): 1–25.
- Puhani, Patrick A. 2012. "The Treatment Effect, the Cross Difference, and the Interaction Term in Nonlinear 'Difference-in-Differences' Models." *Economic Letters* 115: 85–87.
- Wolfers, Justin. 2006. "Did Unilateral Divorce Laws Raise Divorce Rates? A Reconciliation and New Results." *American Economic Review* 96(5): 1802–1820.
- Wu, Lawrence L. 1989. "Issues in Smoothing Empirical Hazards." Pp. 127–59 in Clifford C. Clogg (Ed.), *Sociological Methodology 1989*. Washington: American Sociological Association.

APPENDIX

This appendix provides details on the expression for $dd_{ip(hz)}$ when divorces are governed by the proportional hazard specification in equation (3). Under equation (3), $S(u)$ will be given by:

$$\begin{aligned} S(u) = S(t|t_0) &= \exp\left\{-\int_0^{t-t_0} r_0(s) \exp[b_1 \times g_i + b_2 \times I_i(t) + dd_{hz} \times g_i \times I_i(t)] ds\right\} \\ &= \exp\{-H(t|t_0)\} \end{aligned} \quad (10)$$

where

$$H(t|t_0) = \begin{cases} \exp(b_1 \times g_i)H_0(t|t_0) & t < \tau \\ \exp(b_1 \times g_i)H_0(\tau|t_0) + \exp(b_1 \times g_i + b_2 + dd_{hz} \times g_i)H_0(t|\tau) & t \geq \tau \end{cases} \quad (11)$$

and

$$H_0(b|a) = \int_a^b r_0(s) ds \quad b \geq a, \quad (12)$$

Then plugging equation (12) into equation (2):

$$\begin{aligned} dd_{ip(hz)} &= \{E[p(g=1, I(t)=1)] - E[p(g=1, I(t)=0)]\} - \\ &\quad \{E[p(g=0, I(t)=1)] - E[p(g=0, I(t)=0)]\} \\ &= \{[S_{g=1}(\tau|t_0) - S_{g=1}(\tau_2|t_0)] - [S_{g=1}(\tau_1|t_0) - S_{g=1}(\tau|t_0)]\} - \\ &\quad \{[S_{g=0}(\tau|t_0) - S_{g=0}(\tau_2|t_0)] - [S_{g=0}(\tau_1, |t_0) - S_{g=0}(\tau|(t_0))]\} \\ &= \{2 \times S_{g=1}(\tau|t_0) - S_{g=1}(\tau_1|t_0) - S_{g=1}(\tau_2|t_0)\} - \\ &\quad \{2 \times S_{g=0}(\tau, g=0) - S_{g=0}(\tau_1|t_0) - S_{g=0}(\tau_2|t_0)\} \\ &= \{2 \times S_{g=1}(\tau|t_0) - 2 \times S_{g=0}(\tau|t_0)\} - \\ &\quad \{S_{g=1}(\tau_1|t_0) - S_{g=0}(\tau_1|t_0)\} - \\ &\quad \{S_{g=1}(\tau_2|t_0) - S_{g=0}(\tau_2|t_0)\} \\ &= \{2 \times \exp[-\exp(b_1)H_0(\tau|t_0)] - 2 \times \exp[-H_0(\tau|t_0)]\} - \\ &\quad \{\exp[-\exp(b_1)H_0(\tau_1|t_0)] - \exp[-H_0(\tau_1|t_0)]\} - \\ &\quad \{\exp[-\exp(b_1)H_0(\tau|t_0) - \exp(b_1 + b_2 + dd_{hz})H_0(\tau_2|\tau)] - \end{aligned}$$

$$\begin{aligned}
& \exp[-H_0(\tau|t_0) - \exp(b_2)H_0(\tau_2|\tau)] \} \\
= & \{2 \times \exp[-\exp(b_1)H_0(\tau|t_0)] - 2 \times \exp[-H_0(\tau|t_0)]\} - \\
& \{ \exp[-\exp(b_1)H_0(\tau_1|t_0)] \} - \exp[-H_0(\tau_1|t_0)] - \\
& \{ \exp[-\exp(b_1)H_0(\tau|t_0) - \exp(b_1 + b_2 + dd_{hz})H_0(\tau_2|\tau)] - \\
& \exp[-H_0(\tau|t_0) - \exp(b_2)H_0(\tau_2|\tau)] \} . \tag{13}
\end{aligned}$$

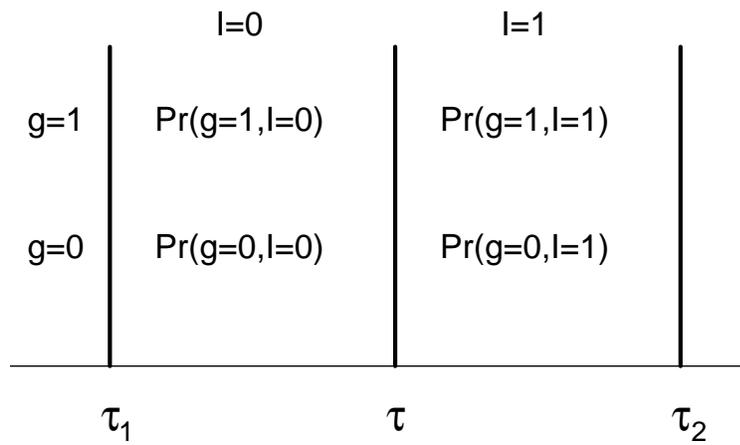


Figure 1: The four components of the linear probability difference-in-differences estimator for two groups and two periods.

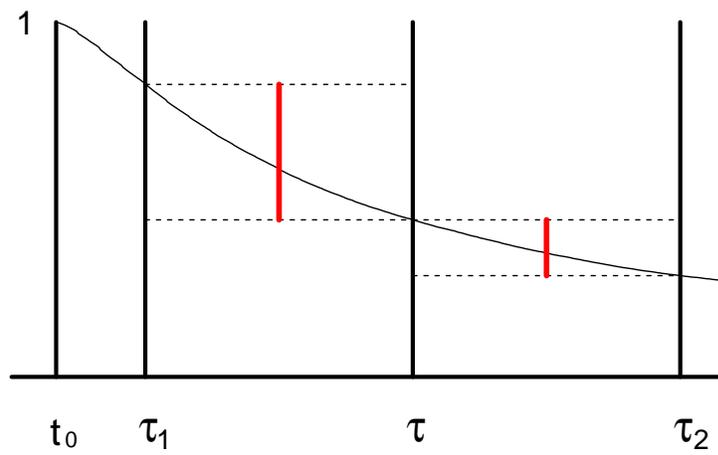


Figure 2: The probability of divorce pre- and post-treatment as a function of the survivor probability when data for divorce are generated by a continuous-time hazard process.

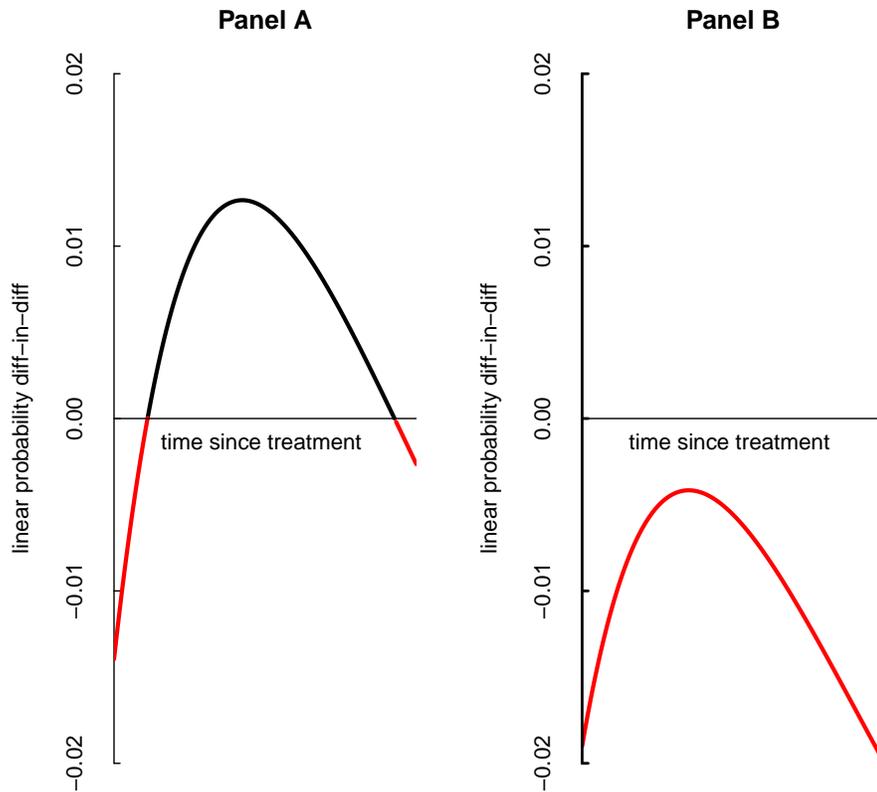


Figure 3: Behavior of $dd_{lp(hz)}$ with time since treatment for an example in which the data are assumed to be generated according to the proportional hazard DD in equation (3). In both Panels A and B, the baseline hazard $r_0(u)$ is set equal to a constant $\lambda = 0.008$; the regression coefficients b_1, b_2 , and dd_{hz} are assumed equal and set to 0.1; the calendar start of marriage, t_0 is set to 0; and the start of observation, τ_1 , is set to 24 months after the start of marriage. In Panel A, treatment begins at $\tau = 60$ months of marital duration, while in Panel B, treatment begins at $\tau = 90$ months of marital duration.

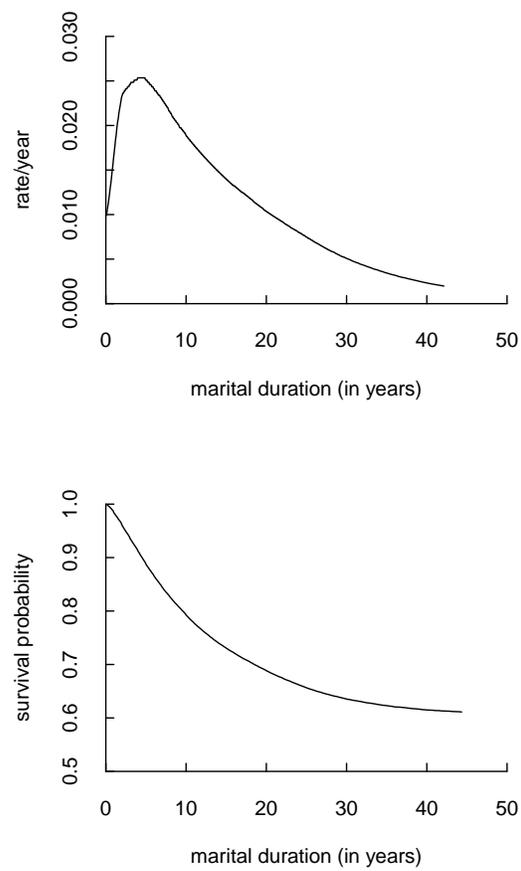


Figure 4: Divorce risks (upper panel) and survivor probabilities (lower panel) by marital duration. U.S. women, June 1980, 1985, 1990, and 1995 Current Population Survey.

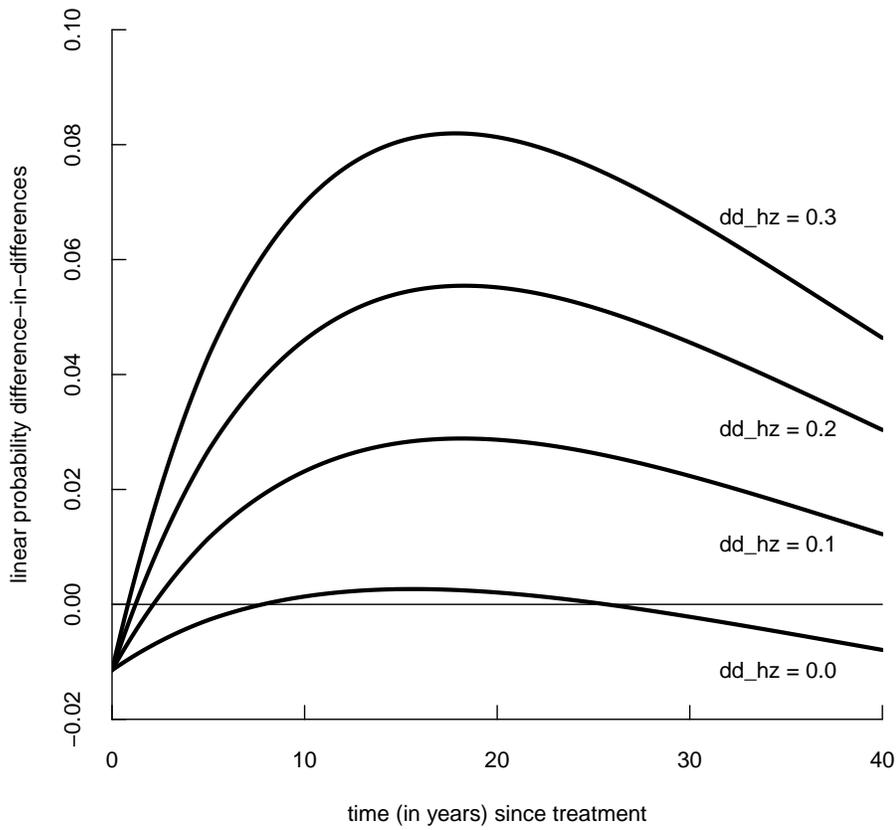


Figure 5: Behavior of $dd_{lp(hz)}$ with time since treatment. Survivor curves for $g = 0$ and $g = 1$. Stylized example assuming $b_1 = b_2 = dd_{hz} = 0.1$ in equation (3) and pre- and post-treatment observation periods of $[60,120]$ and $[120,480]$, respectively. Empirical estimates of baseline divorce risks obtained from retrospective marital histories reported by female respondents in the June 1980, 1985, 1990, and 1995 Current Population Survey.