

Divorce: what does learning have to do with it?*

Ioana Marinescu, University of Chicago

March 14, 2012

Abstract: Learning about marriage quality has been proposed as a key mechanism for explaining how the probability of divorce evolves with marriage duration, and why people often cohabit before getting married. I develop four theoretical models of divorce, three of which include learning. I use data from the Survey of Income and Program Participation to test reduced form implications of these models. The data is inconsistent with models including a substantial amount of learning. On the other hand, the data is consistent with a model without any learning, but where marriage quality changes over time.

*I want to thank Jeff Campbell, Kerwin Charles, Pierre-Andre Chiappori, Analia Schlosser and Rob Shimer for helpful comments on previous versions of this draft. I also want to thank seminar participants at the University of Chicago, CEPR Ammersee 2008, SOLE 2009, NBER Summer Institute 2009 for useful comments.

1 Introduction

Traditionally, the bride and groom leave the parents' nest to be married until death tears them apart. Modern marriages fundamentally differ from this model in that the rate of divorce is substantial. This raises the key question of why people marry only to divorce a few years later. The answer must be that something changes over time that makes divorce preferable to the continuation of the relationship. A number of theories exist that posit different sources of change as being the drivers of divorce. One theory is that the outside option of one of the partners can change, and the marriage dissolves when one of the partners meets a better match (Weiss and Willis, 1997). A second theory is that one of the aims of marriage is consumption insurance and marriages end when such insurance fails (Hess, 2004). Finally, a third theory invokes learning: when spouses marry they are not perfectly informed about their match quality. Instead, they learn about it over time, and divorce occurs when spouses find out that they are in a bad match. This theory also offers a reason for why people cohabit before marriage, another key feature of modern relationships (e.g. Brien, Lee, Lillard and Stern, 2006). In this framework, cohabitation allows partners to learn about match quality before making a commitment that is costly to break.

I argue that the theories explaining why people marry only to divorce a few years later can be broadly classified into two categories. One theory is that the value of the relationship relative to the outside option changes over time. The other is that the value of the relationship does not change over time, but partners' beliefs about this value change as they gradually learn about match quality. This paper will show that, empirically, learning plays little to no role in divorce and that divorce can be fully explained by changes in the value of the relationship itself. I start with constructing a model that nests both theories by assuming that match quality follows a random walk and partners learn about match quality through signals a la Jovanovic (1979). I then derive empirically testable predictions that can, under some parameter restrictions, distinguish between a

pure learning model (no changes in match quality over time), a pure changes model (no learning) and a mixed model that includes both learning and changes in match quality. In order for models to yield starkly different predictions, it is necessary to assume that there is substantial learning in any of the models that includes learning, i.e. that the signal of marriage quality observed by couples is noisy enough. Under this assumption and a number of additional parameter restrictions, we can make the following empirically testable predictions. First, the divorce hazard increases and then decreases with marriage duration in the pure learning model or the mixed model, but it monotonically decreases with duration in the pure changes model. In order to get further predictions, I derive the impact of observing a signal of low match quality on the divorce hazard. I find that, in the pure learning model, marriages for which a signal of low match quality is observed are more likely to terminate even many periods after the signal of low match quality was observed. In the other models, the impact of observing a signal of low match quality declines to zero as time goes by: after enough time has elapsed, marriages for which a signal of low match quality was observed at some time in the past are no more likely to terminate than marriages for which no such signal was observed. A second prediction using a signal of low match quality is that in the pure learning model, the impact of a signal of low match quality on the divorce hazard monotonically decreases with marriage duration. In contrast, in the pure changes model, the impact of a signal of low match quality on the divorce hazard monotonically increases with marriage duration. Finally, in the mixed model, the impact of a signal of low match quality on the divorce hazard decreases and then increases with marriage duration.

My empirical analysis uses monthly longitudinal data on married and cohabiting couples from the 1990-2004 waves of the Survey of Income and Program Participation (SIPP). I use job loss (either being laid off or getting fired) as a signal of low match quality. Indeed, spouses care about economic success (e.g. Hitsch et al., 2010). Additionally, job loss has a negative impact on subsequent earnings, unemployment (Gibbons and Katz, 1991) and survival (Sullivan and von Wachter, 2009). Importantly, job loss is in

fact associated with an increased probability of divorce (Charles and Stephens, 2004). Using the tests outlined above, I find that the data is most consistent with the pure changes model. Substantively, I find that the divorce hazard monotonically decreases with marriage duration and that the impact of a discharge for cause on the divorce hazard is monotonically increasing with duration¹. Additionally, I find that a job loss that occurred more than a year ago has no significant impact on current divorce. One may question the assumption that job loss is simply a signal of low match quality: what if job loss has a causal effect and actually decreases match quality just as much for good as for bad marriages? I use an alternative learning model that embodies this assumption. I cannot reject that this model is incorrect. Since I can reject the predictions of models with a large amount of learning, and I cannot reject that the predictions of the pure learning model are incorrect, I infer that learning does not play an important role in explaining divorce. By contrast, if I assume that the pure changes model makes incorrect predictions, I can reject this hypothesis. I conclude that the pure changes model is the best candidate to explain the divorce hazard.

This paper makes three key contributions. First, while the learning model has been widely used to explain divorce and cohabitation, I show that learning plays at best a modest role in accounting for how divorce probabilities change with marriage duration. On the other hand, a model that assumes that match quality is perfectly observed and follows a random walk can fully explain the data. This finding can also make sense of the fact that it has been difficult to show empirically that cohabitation decreases the divorce hazard, a key implication of the learning model. Beyond the fact that there is selection into cohabitation, learning may simply not play an important role in marriage. The second contribution of this paper is to the literature on the impact of job loss on divorce. While it is known that job loss is associated with a higher divorce hazard, it has not yet been clear to what extent this relationship could be interpreted causally. The results of

¹As will be discussed below, it is likely that other studies have underestimated the divorce hazard at low durations because no high frequency panel data was available. Underestimating the hazard at low durations will tend to yield an overall hazard that increases initially with relationship duration.

this paper are consistent with job loss having a causal impact on divorce. For a given belief about match quality prior to job loss, marriages in which a job loss occurs are more likely to end in divorce. At the same time, job loss does not occur randomly: instead, the evidence is consistent with lower quality marriages being also more likely to experience job loss, and in particular a discharge for cause. The third contribution of this paper is to provide a model of relationship separation that can be applied to other relationships such as employment relationships or commercial contracts. The model yields empirically testable predictions that can allow us to learn about the role of learning and shocks in other domains.

There is a limited literature in economics that investigates the impact of labor market shocks on the probability of divorce or separation. Weiss and Willis (1997) look at the impact of unexpected wage gains on divorce and find a negative impact for men's wage gains and a positive one for women's wage gains. This supports the idea that women prefer men with higher earning potential. Charles and Stephens (2004) find that the probability of divorce increases when either spouse is laid off (with a stronger effect for men). Moreover, a layoff has a stronger effect than a plant closing. Charles and Stephens (2004) speculate that what matters is the information conveyed by job loss about the fitness of the partner as a mate rather than purely economic factors. Plant closure also significantly increases the probability of divorce (Rege, Telle and Votruba, 2007), which suggests that job loss has a causal impact on divorce.

There is a much larger literature in psychology that addresses marital functioning and its relationship to economic factors, even though this literature does not focus specifically on the impact of job loss. Economic stress decreases marital satisfaction, and this is in part due to worse marital functioning, i.e. worse communication and the like (Conger et al. 1999). Kinnunen and Feldt (2004) show that even in a country like Finland, where unemployment benefits are very generous, the longer the husband stays unemployed the more likely his wife is to report increased conflict and decreased common interests.

This paper also relates to a theoretical literature that explains the evolution of the

hazard of relationship separation in various contexts. With respect to job separation, Jovanovic (1979) develops the classical learning model and Mortensen and Pissarides (1994) explain job separation through the occurrence of random productivity shocks. A series of subsequent papers have further developed theory and tested it in the context of job separation (Farber, 1994, Nagypal, 2007, Marinescu, 2009, Kahn and Lange, 2010). In the case of marital separation, Brien, Lee, Lillard and Stern (2006) develop a marriage model inspired by Jovanovic (1979) and structurally estimate the model, finding that cohabitation is explained by both the need to learn about potential partners and by the desire to hedge against future bad shocks. Finally, an extension of the Jovanovic (1979) learning model has also been developed and tested for firm learning in the first year of a firm's life (Abbring and Campbell, 2005): there is no evidence for Jovanovic-style learning in this context. With the exception of Farber (1994), all of the papers that test theory empirically use structural estimation. This paper shows how one can use intuitive and easy to implement reduced-form tests to test for the presence of substantial learning. Additionally, to perform these tests, one only needs a crude signal of low match quality: a dummy variable is sufficient. More detailed data is of course in principle desirable but it will tend to be missing for some applications, and in particular in the case of marriage. Indeed, there is to my knowledge no high-frequency data that tracks beliefs about match quality in marriage.

Finally, this paper relates to a small econometrics literature that addresses the issue of causality in duration models. The question is whether some treatment is causally related to duration, or whether the impact of the treatment is due to unobserved heterogeneity. Abbring and van den Berg (2003) develop empirical tests that are similar in spirit to what I propose here. A contribution of my paper is to show how these econometric tests can be grounded in microeconomic theory based on agents' optimizing behavior.

The remainder of the paper is organized as follows. Section 2 presents a theory of marriage duration. Section 3 discusses the main empirical results. Section 4 discusses some further empirical results and robustness tests, and section 5 concludes.

2 Theory

2.1 Model specification

The model analyzes the decision of one of the partners in a relationship to continue that relationship or end it. To make the discussion more intuitive, I will assume that the woman is the partner who decides whether to separate. The problem is of course symmetric for the man². I assume that individuals are matched randomly, which implicitly assumes that there are search frictions. The woman entering a new marriage does not know the exact value of such a marriage. The quality of the marriage, or match quality, is what makes the marriage valuable to the woman: monetary benefits, companionship, love, children, etc. Thus, the quality of a marriage is multi-dimensional. To simplify the modeling, I assume that this multi-dimensional marriage quality is reducible to a single index. The woman holds a prior belief about the distribution of quality in the population of potential relationships (partners). This prior $P(q_0)$ is normally distributed with mean \bar{q} and variance σ_0^2 . Denote by q_k the (true) match quality at length k . As long as the marriage continues, match quality is assumed to evolve over time according to the following AR(1) process:

$$q_k = \rho q_{k-1} + c + \epsilon_{k-1}^q \quad (1)$$

where $\epsilon_k^q \sim N(0, \sigma_p)$. c is a deterministic drift. For simplicity, I further assume that $\rho = 1$ and $c = 0$, so that the process is a random walk³. At each period, the woman observes a noisy signal of the quality of the partner. The signal of match quality is an observation z_k defined as:

$$z_k = q_k + \epsilon_k^z \quad (2)$$

where $\epsilon_k^z \sim N(0, \sigma_{obs})$. The best estimate \hat{q}_k of q_k given all observations $z_1 \dots z_k$ is fully determined by the Kalman filter solutions (see Arulampalam et al. (2001)).

²If utility is transferable, then the woman and the man will always agree on the separation, so it does not matter which point of view we consider.

³Assuming some slightly smaller ρ and higher c does not affect the qualitative results.

Figure 1 shows the timing of the woman's decision process. At every time step⁴, the woman has two possible actions a : she can continue the current marriage ($a = C$) or separate from the current partner, pay a separation cost f , and begin a new marriage with another partner⁵ ($a = S$). The woman chooses one of these actions depending on her current belief and the length of the marriage k . Define a policy π , which gives for each belief and marriage length the action to be taken. Define the Q function $Q^\pi(\hat{q}_k, a)$ as the expected return of taking action a today and then following the policy π in the future. The value function $V^\pi(\hat{q}_k)$ gives the current and future rewards of the woman as a function of current belief, assuming that the woman follows policy π from now on. The optimal policy maximizes $V^\pi(\hat{q}_k)$, and gives rise to the optimal value function $V^*(\hat{q}_k)$. The optimal action value function Q^* is defined as a function of the optimal value $V^*(\hat{q}_k)$:

$$Q^*(\hat{q}_k, C) = \hat{q}_k + \delta \sum_{\hat{q}_{k+1}} P(\hat{q}_{k+1}|\hat{q}_k) V^*(\hat{q}_{k+1}) \quad (3)$$

$$\begin{aligned} Q^*(\hat{q}_k, S) &= \bar{q} - f + \delta \sum_{\hat{q}_1} P(\hat{q}_1|\bar{q}) V^*(\hat{q}_1) \\ &= V_{new} - f, \text{ where } V_{new} = \bar{q} + \delta \sum_{\hat{q}_1} P(\hat{q}_1|\bar{q}) V^*(\hat{q}_1) \end{aligned} \quad (4)$$

The optimal value is given by the Bellman equation:

$$V^*(\hat{q}_k) = \max_{a \in [C, S]} Q^*(\hat{q}_k, a) \quad (5)$$

Thus, if the woman decides to continue with the current relationship, she gets the expected value of the marriage quality in the current period plus any future rewards. If she instead decides to end the marriage, she gets the average value of a new match randomly drawn from the prior distribution plus any future rewards from that new match⁶, and

⁴I assume that the length of the marriage is limited to some length k_{max} .

⁵The new marriage begins instantly and no one remains unmatched. The costs of switching partners are captured by f in this model.

⁶The definition of the reward function is compatible with a Nash bargaining solution where the two partners split the surplus, so that, while the marriage continues, each partner gets a fixed share. Suppose that $0 \leq \alpha \leq 1$ is the share of the rewards from the marriage received by the woman. The reward of the woman would then be $\alpha \hat{q}_k$ if continuing and $\alpha \bar{q}$ if separating; but this change is not substantial since it

has to pay f . The cost f could depend on marriage duration k but, to keep the model parsimonious, I abstract from such a dependency here ⁷. The separation cost f covers the direct cost of ending the current relationship, such as a paying for a divorce lawyer. It also covers the costs of beginning a new marriage, such as fees to enroll in online dating sites.

In this framework, the optimal policy followed by the woman is uniquely defined by $\tau(k)$, the belief such that the woman is indifferent between continuing and separating from her partner at marriage length k . In other words, the threshold for separation $\tau(k)$ is defined by the equalization of Q functions for the actions “continue” (equation (3)) and “separate” (equation (4)). To compute the optimal values and policy, one fixes parameter values and uses a version of the “value iteration” algorithm⁸, which has been shown to converge to the solution of a Partially Observed Markov Decision Problem such as the one we have here (see Hauskrecht(2002)).

Note that the planning horizon of the woman is assumed to be infinite. This means that the woman lives indefinitely; or alternatively, the woman’s retirement from the marriage market is at some date so far away in the future that given the discount factor, it does not play any role in the woman’s current decisions. The model may thus not be adequate for explaining the behavior of older women.

The model developed here embeds two polar cases. If $\sigma_p = 0$, match quality is constant over the duration of the marriage. In that case, and assuming further that $\sigma_{obs} > 0$, the model is essentially identical to Jovanovic (1979): it is a model of learning about match quality, or “pure learning model”. If $\sigma_{obs} = 0$ and $\sigma_p > 0$, match quality is perfectly observed but evolves over time. One example of such a model is Mortensen and Pissarides (1994), where job destruction is generated by idiosyncratic shocks to job

simply amounts to rescaling the distribution of match quality.

⁷In fact, one can assume any deterministic relationship between f and k without further technical complications in solving the model.

⁸Briefly, one starts with an arbitrary value for V_{new}, V_{new}^0 . At the last period of the marriage, the marriage ends, and so the value of the marriage is $V_{new}^0 - f$. One then computes the value of the marriage at the period prior to the last period, and so on recursively until the first period. One gets a value for the marriage at the first period V_{new}^1 . One starts the recursion again with $V_{new} = V_{new}^1$ at the last period of the marriage, and one does this until $V_{new}^n \approx V_{new}^{n-1}$.

productivity. In the context of a couple’s separation, if $\sigma_{obs} = 0$ and $\sigma_p > 0$, then separation is only due to changes in marriage quality. This is a “pure changes model”. When $\sigma_{obs} > 0$ and $\sigma_p > 0$, this is a “mixed model” with both learning and changes to marriage quality, i.e. the woman learns from noisy observations about a marriage quality that is constantly evolving. Because the small literature on marriage duration (e.g. Svarer, 2004) proposes the pure learning model as an explanation of the probability of marriage dissolution as a function of marriage length, this paper takes that model as a starting point. The question is then whether one can do better in explaining empirical data about marriage duration by using changes in marriage quality.

2.2 Empirically observable outcomes

I discuss three empirically observable outcomes that can be derived from the model. I will then discuss how these outcomes differ depending on model parameters.

2.2.1 Divorce hazard

The theoretical hazard of separation is the result of infinitely many women confronted with the same separation decision problem; it summarizes the average separation behavior of women over marriage lengths. One can compute the theoretical separation hazard once the threshold for separation is known. Note that at marriage length zero, when no observation has been made yet, $\hat{q}_0 = \bar{q}$ for all matches. For all women, the belief is the same as the prior. Let $p_k(\hat{q}_k)$ be the density of women who hold a belief with mean \hat{q}_k at length k , given that they follow the optimal policy embodied in $\tau(k)$. The hazard of separation at length k , $h(k)$, can be computed recursively, starting at $k = 1$ and using the distributional assumptions. The initial values for the distribution of women’s expected

beliefs about match quality are:

$$p_0(\hat{q}_0) = \begin{cases} 1 & \text{if } \hat{q}_0 = \bar{q} \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

$$p_1(\hat{q}_1) = \sum_{\hat{q}_0} p_0(\hat{q}_0) P(\hat{q}_1 | \hat{q}_0) = P(\hat{q}_1 | \bar{q}) \quad (7)$$

The hazard of separation at length k , $h(k)$, can then be computed recursively, starting at $k = 1$:

$$h(k) = \sum_{\hat{q}_k = \hat{q}_{min}}^{\hat{q}_k = \tau(k)} p_k(\hat{q}_k) \quad (8)$$

$$p_k(\hat{q}_k) = 0 \text{ if } \hat{q}_k \leq \tau(k) \quad (9)$$

$$p_k(\hat{q}_k) = \frac{p_k(\hat{q}_k)}{\sum p_k(\hat{q}_k)} \quad (10)$$

$$p_{k+1}(\hat{q}_{k+1}) = \sum_{\hat{q}_k} p_k(\hat{q}_k) P(\hat{q}_{k+1} | \hat{q}_k) \quad (11)$$

Equation (10) insures that the mass of women is always normalized to 1. $P(\hat{q}_{k+1} | \hat{q}_k)$, the distribution of possible beliefs at marriage duration k given the belief at $k - 1$, can be computed given distributional assumptions. Once the divorce hazard is computed, one can analyze how it changes with marriage duration. It is to be expected that different model parameters will yield different shapes for the divorce hazard.

2.2.2 Impact of job loss on divorce hazard as a function of marriage duration

Empirically, the divorce hazard is the key observable outcome. But there is only so much information one can extract from a single divorce hazard. Further information can be provided by examining the divorce hazards of population subgroups that differ in ways that are interpretable within the model. Specifically, we can analyze the divorce hazard for couples that received a negative signal about marriage quality versus couples that did not receive such a negative signal. We can further analyze how the impact of a negative

signal about marriage quality changes with marriage duration.

To use a concrete and empirically observable case, assume that a job loss corresponds to the woman observing a negative signal of marriage quality. The fact that a husband's losing a job is bad news is empirically plausible given the importance of financial considerations for mate choice and marriage stability. However, losing a job is just a signal because the wife cannot perfectly observe how much the husband is to blame for the job loss. If the husband has little responsibility in the job loss, his future economic prospects are less likely to be negatively affected. To model this situation, I assume that if the man experiences a job loss⁹ at marriage length k , then the woman observes $z_k < z^*$, where z^* is some low marriage quality threshold (this threshold has to be lower than the mean of the prior, i.e. the average quality of a new marriage). The idea is that a job loss is a negative signal about marital quality, but I do not take a stance about exactly how bad it is; I just assume that it is below some threshold. Since \hat{q}_{k-1} and z_k uniquely determine \hat{q}_k given distributional assumptions, I define a function $g(\hat{q}_{k-1}, z^*)$ that gives, for each \hat{q}_{k-1} , the value of \hat{q}_k corresponding to the observation of z^* at period k . Since for a given \hat{q}_{k-1} , \hat{q}_k increases in z_k , $z_k < z^*$ can be rewritten as $\hat{q}_k < g(\hat{q}_{k-1}, z^*)$. The probability that a match of quality \hat{q}_{k-1} is dissolved at k given $z_k < z^*$ is then:

$$P(\hat{q}_k < \tau(k) | \hat{q}_{k-1}, \hat{q}_k < g(\hat{q}_{k-1}, z^*))$$

Using the definition of conditional probabilities and Bayes' rule:

$$P(\hat{q}_k < \tau(k) | \hat{q}_{k-1}, \hat{q}_k < g(\hat{q}_{k-1}, z^*)) = \frac{P(\hat{q}_k < \min(g(\hat{q}_{k-1}, z^*), \tau(k)) | \hat{q}_{k-1})}{P(\hat{q}_k < g(\hat{q}_{k-1}, z^*) | \hat{q}_{k-1})}$$

It is important to realize two things with respect to the equation above. First, I assume

⁹Given the assumptions of the model, a negative observation about the current marriage does not affect the value of alternative marriages. When it comes to job loss, this assumption should hold in that the husband's job loss should not affect the wife's prospects on the marriage market. However, the husband's job loss is likely to affect his prospects on the marriage market. In that sense, in order for the reasoning presented here to fully carry through, it must be the case that the husband cannot compensate the wife for the decline in the value of the marriage.

that job loss ($z_k < z^*$) happens at period k and make no specific assumption about whether job loss occurred earlier than k . The history of observations up to z_{k-1} is simply assumed to be consistent with the distributional assumptions and the optimal strategy of the woman. Second, the denominator of the equation above shows that the probability of job loss at period k depends on estimated marriage quality at period $k - 1$. This allows for job loss to occur more often in low quality marriages: since $z_k = q_k + \epsilon_k^z$ (see equation 2), it is straightforward to see that observing $z_k < z^*$ is more likely if q_k is low, and hence if \hat{q}_k is low.

The hazard h_b given a bad observation $z_k < z^*$ can then be written as:

$$h_b(k) = \sum_{\hat{q}_{k-1}} p_{k-1}(\hat{q}_{k-1}) \frac{\sum_{\hat{q}_k = \hat{q}_{min}}^{\hat{q}_k = \min(g(\hat{q}_{k-1}, z^*), \tau(k))} P(\hat{q}_k | \hat{q}_{k-1})}{\sum_{\hat{q}_k = \hat{q}_{min}}^{\hat{q}_k = g(\hat{q}_{k-1}, z^*)} P(\hat{q}_k | \hat{q}_{k-1})} \quad (12)$$

Similarly, the hazard h_g given a relatively good observation $z_k > z^*$ is:

$$h_g(k) = \sum_{\hat{q}_{k-1}} I[g(\hat{q}_{k-1}, z^*) \leq \tau(k)] p_{k-1}(\hat{q}_{k-1}) \frac{\sum_{\hat{q}_k = g(\hat{q}_{k-1}, z^*)}^{\hat{q}_k = \tau(k)} P(\hat{q}_k | \hat{q}_{k-1})}{\sum_{\hat{q}_k = g(\hat{q}_{k-1}, z^*)}^{\hat{q}_k = \hat{q}_{max}} P(\hat{q}_k | \hat{q}_{k-1})} \quad (13)$$

where I is an indicator function. The indicator function I is necessary because the hazard is positive only if $g(\hat{q}_{k-1}, z^*) \leq \tau(k)$; otherwise not observing a job loss guarantees that no divorce occurs at k .

Empirically we can estimate $\log[h_b(k)/h_g(k)]$, i.e. the impact of job loss on the divorce hazard, and we can examine how this impact varies with marriage duration.

2.2.3 Impact of job loss on the divorce hazard as a function of time elapsed since job loss

So far, I have examined the effect of observing a job loss at length k on separation at length k . However, it is also interesting to ask how this effect evolves over time: relative to those marriages that did not experience a job loss at length k , how much more likely are marriages that did experience a job loss at length k to dissolve at lengths $k + 1$, $k + 2$,

..., $k+n$? One can answer this question by computing the hazard of separation at lengths $k+1, \dots, k+n$ separately for those marriages that did experience a job loss at k and those that did not. I can then analyze how the impact of job loss varies with time elapsed since job loss.

2.2.4 Causality, learning and job loss

When estimating econometric models, one often asks if some variable is “causally” related to the outcome of interest. In this specific case, one may ask whether job loss is causal in generating divorce. In the models discussed above, job loss is causal in the sense that observing job loss at the current period increases the divorce hazard for a given belief at the previous period. However, it is important to point out that in the models discussed above, the impact of job loss on divorce is *also* due to selection in the sense that job loss is more likely to happen in bad marriages. Indeed, since I formalize job loss as $z_k < z^*$, and $z_k = q_k + \epsilon_k^z$, then clearly the probability of observing $z_k < z^*$ is higher when actual match quality q_k is lower. In that sense, job loss in these models does not have a purely causal effect. Job loss happens more often to worse quality matches, and it further causes them to dissolve.

Since the impact of job loss on divorce in the models discussed above is not purely causal, it is useful to formulate an alternative model where the impact of job loss is purely causal, i.e. where job loss is not correlated with match quality. I start with the model above in the pure learning case ($\sigma_p = 0$) and add some shocks that are independent of marriage quality. Call this model “learning with shocks”. Assume that there is a two-states Markov process and the marriage can be either in the “job loss” state, or in the “no job loss” state. There are two transition probabilities, p_1 and p_2 , that govern the transitions to and from the “job loss” state. Importantly, these transition probabilities are independent of match quality. Assume that the optimal action value function if continue $Q^*(\hat{q}_k, C)$ is defined as a function of the optimal value $V^*(\hat{q}_k, j)$, where j identifies whether

the match is in the “no job loss” ($j = 1$) or “job loss” ($j = 2$) state:

$$Q^*(\hat{q}_k, 1, C) = \hat{q}_k + \delta \sum_{\hat{q}_{k+1}} P(\hat{q}_{k+1}|\hat{q}_k)(1 - p_1)V^*(\hat{q}_{k+1}, 1) + P(\hat{q}_{k+1}|\hat{q}_k)p_1V^*(\hat{q}_{k+1}, 2) \quad (14)$$

$$Q^*(\hat{q}_k, 2, C) = \hat{q}_k - b + \delta \sum_{\hat{q}_{k+1}} P(\hat{q}_{k+1}|\hat{q}_k)(1 - p_2)V^*(\hat{q}_{k+1}, 2) + P(\hat{q}_{k+1}|\hat{q}_k)p_2V^*(\hat{q}_{k+1}, 1) \quad (15)$$

where $b > 0$ is a constant that embodies the per period cost of job loss.

Assume further that all marriages start in the “no job loss” state. Thus, the optimal action value function if separating is the same as in the pure learning model. This model can be solved numerically just as the models discussed above. Note that it is similar in spirit to the mixed model since it includes both learning and changes in the value of the marriage over time. However, it differs from the mixed model in that job loss is a transitory negative shock to the rewards from marriage that is *independent* of match quality.

2.3 Distinguishing between models on the basis of empirically observable outcomes

2.3.1 Calibration: reference parameters

I start with a reference case for the parameters. When empirical evidence is available, I calibrate the parameters, and I otherwise choose reasonable values using considerations specific to each parameter.

I use up to 50 periods of marriage, and the period is to be understood as a year. Consequently, I choose a value of 0.8 for the discount factor, consistent with empirical evidence about the yearly subjective discount factor (Frederick et al., 2002). The cost of divorce is best understood relative to the per period reward of an average new marriage (mean of prior). If we greatly simplify and assume that the benefits and costs of marriage are mostly financial, the empirical literature suggests that women typically experience

a 30% to 40% adjusted income drop in the aftermath of divorce (Smock, Manning, and Gupta 1999, Gadalla 2008). Gadalla (2008) shows that the income drop persists for at least 5 years post divorce. If we assume a 35% income drop spread over 5 years and discounted using the 0.8 discount factor, this amounts to a present value of 94% of pre-divorce income, i.e. close to the income of a year of marriage pre-divorce. Using this calculation, it seems that a value of 30 is reasonable since it assumes that the cost of divorce is equal to the value of a year of a new marriage. It is likely that a value of 30 for the cost of divorce underestimates the total cost. Indeed, the calculation only takes into account income, but other costs exist such as health costs (Amato, 2010), and they are quite difficult to quantify. On the other hand, the financial cost of divorce for men is lower than for women (Gadalla, 2008). While men are less likely to initiate divorce, some do, and therefore a value of 30 for the divorce cost represents a reasonable compromise between the costs of divorce for women and men.

The reference values for σ_{obs} and σ_p were chosen with the view that they should represent substantial learning and substantial changes in match quality. Given that the standard deviation of the prior is 5, how large these parameters are needs to be gauged in reference to that value. I chose $\sigma_{obs} = 10$ for the reference case, which represents substantial learning in the sense that the signal has a standard deviation that is twice as high as the standard deviation of the prior. For the models with changes in relationship quality, I chose $\sigma_p = 5$, which is indeed a high value as the standard deviation of shocks to match quality is as large as the standard deviation of the prior. Values that are much larger than this are not empirically plausible. Indeed, if the standard deviation of the shocks to match quality is much larger than that of the prior, then your average new husband has a non negligible chance of becoming worse in the second year of marriage than the very worst husband was in the first year of marriage. Such a rapid change does not seem plausible.

For the learning with shocks model, I assume that being in the job loss state takes 10 units off the quality of the marriage, which is a 33% decline at the mean of the prior.

Given the high costs of job loss, such a magnitude seem plausible. Then, in keeping with the empirical data described below, I choose transition parameters implying that job loss is fairly uncommon ($p_1 = 0.05$). I further assume that the job loss state is persistent ($p_2 = 0.1$), consistent with the empirical literature showing the long-term negative consequences of job loss (Gibbons and Katz, 1991, Sullivan and von Wachter, 2009).

These reference parameters being fixed (see Table 1 for the full list), we can analyze how different model parameters impact empirically observable outcomes.

2.3.2 Outcomes in the reference case

As is well known since Jovanovic (1979), the pure learning model generates a divorce hazard that increases and then decreases with marriage duration (see Figure 2). The initial increase is due to the fact that initially match quality is not well known and, as long as separation costs are not zero, it is worth waiting to gain more information. As sufficient information accumulates, more and more of the relatively bad matches dissolve, which increases the separation hazard. Eventually, the hazard decreases as continuing marriages are dominated by higher quality matches. By contrast, the pure changes model generates a hazard that decreases monotonically with marriage duration. Like the pure learning model, the mixed model and the learning with shock model generate a hazard that increases and then decreases with marriage duration (Figure 2).

Figure 3 plots the impact of job loss on the divorce hazard (log hazard ratio $\log[h_b(k)/h_g(k)]$) under the various models under consideration. In the pure learning model, the log hazard ratio decreases very rapidly and then more slowly as marriage duration increases. Observing a job loss early on in the marriage is very informative about marital quality, but becomes less informative as more information is gathered. In the pure changes model, the log hazard ratio increases with marriage duration. The intuition for the increase in the hazard ratio is as follows. First note that, in the pure changes model, the impact of job loss relative to no job loss is largest for “mediocre” marriages. Very bad marriages

are quite likely to end regardless of whether or not a job loss occurs: since marital quality follows a random walk, a marriage that is in the vicinity of the divorce threshold is likely to be dissolved soon. Similarly, for marriages that are really far away from the threshold, divorce is unlikely whether or not job loss occurs. The reason why the log hazard ratio increases over time in the pure changes model is that the proportion of mediocre marriages increases. At low marriage durations, the distribution of estimated marital quality is roughly a truncated normal. Truncation is driven by the divorce threshold and occurs to the left of the mode of the distribution: because there are positive divorce costs, the threshold for separation is lower than mean marriage quality for new partners. At low marriage durations, the bulk of marriages are therefore far away from the threshold. However, as shocks are added to marital quality, the distribution becomes more spread out, and so the proportion of marriages that are “mediocre” increases. This increase in the proportion of mediocre marriages increases the risk of divorce in the case where the marriage is affected by job loss. Therefore, the impact of job loss on divorce increases with marriage duration. In the mixed model, the log hazard ratio first decreases with marriage duration, then subsequently increases again. The initial decrease is due to learning ($\sigma_{obs} > 0$) while the subsequent increase is due to changes in marriage quality ($\sigma_p > 0$). To check for the roles of learning versus changes in match quality, one can vary σ_{obs} and σ_p in the mixed model. Thus, one can see that the initial decrease in the hazard ratio is stronger with a higher σ_{obs} , i.e. when there is more scope for the woman to keep learning new things about the marriage as time goes by. The subsequent increase in the hazard ratio is stronger as σ_p increases¹⁰. Finally, Figure 3 shows that the log hazard ratio decreases and then increases with marriage duration in the learning with shocks model, a prediction which is similar to the prediction for the mixed model.

Figure 4 plots the log ratio of the hazard if a bad signal was observed at length 5 (i.e. if job loss occurred), to the hazard if no bad signal was observed at length 5. This is not

¹⁰An increase in σ_p increases both h_b and h_g . However, while σ_p increases h_b by roughly the same amount at all marriage durations, σ_p increases h_g much more at low marriage durations. This is what drives the steeper increase in the hazard ratio h_b/h_g when σ_p increases.

defined in the pure changes model as formulated above. Since the signal is a perfectly accurate reflection of true marriage quality, those marriages that get a bad signal at tenure k dissolve immediately if the signal is below the threshold, and so there are no marriages left to be dissolved at $k + 1$ or later. To get a more interesting prediction, I allow the signal to be noisy, i.e. instead of representing job loss by $z_k < z^*$, I represent it by $z_k < z^* + \epsilon_k$, where ϵ_k is independently normally distributed with mean 0 and a standard deviation of 5. This adds a lot to the computational burden since one must loop over a discretized list of values for ϵ_k . On the other hand, there is little gain to this extension in terms of qualitative results¹¹. For this reason, I only use this extension of the pure changes model here. I call this model “pure changes model with noise”.

In Figure 4, the log hazard ratio is plotted against time elapsed since the job loss occurred. The figure shows that in all models the hazard ratio is largest in the periods immediately following period 5. The mixed model and the pure changes model yield very similar results. The key difference between the pure learning model and the mixed model or the pure changes model is how the log hazard ratio evolves over time. In the pure learning model, the hazard for those marriages that got a bad observation at 5 remains larger than the hazard for those marriages that did not get a bad observation at 5. In contrast, in the mixed model or the pure changes with noise model, the difference between the two hazards decreases quickly with time. These qualitative conclusions do not depend on the choice of $k = 5$. The intuition for these results is as follows. In the pure learning model, the true quality of the marriage is fixed. As a result, marriages that get a bad observation at 5 are on average worse than those that do not, and this difference persists over time. For both those marriages that get a bad observation at 5 and the others, learning continues until hazards converge to zero. However, until then, there will be more separations among marriages that get a bad observation at 5 because they are worse on average. In the mixed model or the pure changes with noise model, marriage quality changes continuously. As in the pure learning model, those marriages that get a

¹¹I verified that the qualitative results in cases other than the pure changes model are unchanged for a range of parameters.

bad observation at 5 are on average worse than those that do not. However, since true marriage quality evolves according to a random walk after length 5, their quality at length 5 is less and less informative about their present quality as time goes by, and after some time the two groups no longer differ in their separation hazards¹². In the learning with shocks model, the log hazard ratio decreases with time elapsed since job loss as for all the other models. However, what is interesting here is that the log hazard ratio becomes smaller than zero when enough time has elapsed since job loss occurred (see Figure 4). In other terms, when enough time has elapsed since job loss, the hazard of divorce for those couples that experienced a job loss at k is *smaller* than the hazard of divorce for those that did not experience such a job loss at k . This is quite intuitive: since job loss is transitory and independent of match quality in the pure learning with shocks model, those couples that survive after experiencing job loss have on average a higher match quality than those that did not experience a job loss. This prediction differentiates this model from all the models I examined above.

Table 2 summarizes empirical predictions that can allow us to distinguish between alternative models, assuming that parameters are as in the reference case.

2.3.3 Sensitivity to changes in reference parameters

To assess how robust differences between models are to changes in parameters, I have computed model outcomes for a large number of alternative parameter combinations. Table 3 lists the parameter ranges I have explored. For the standard deviation of the observation and of the process, I have chosen a fairly small step starting from 0, so that we can see how predictions change starting from no learning or no changes in match quality. For the standard deviation of the observation, I chose twice the reference case as the upper bound. For the standard deviation of the process, I have also chosen twice the reference case as the upper bound. For the divorce cost, I have explored a slightly lower value than in the reference case (I did not look at even smaller values because such

¹²This also holds if the AR(1) process is stationary. Indeed, since all marriages converge to the long-run mean, past history becomes less and less relevant.

a low divorce cost is inconsistent with empirical data), and also several higher values, up to twice the divorce cost in the reference case. Finally, for the threshold for a bad observation, I have chosen a broad range of values within the range 0 to 30; indeed, since a bad observation must imply that the relationship is below average, this threshold has to be below 30.

Table 4 shows which parameter ranges (within the range of parameters explored in Table 3) preserve the same predictions as in the reference case. For each model, outcome and parameter, the table shows which parameter range preserves the reference case prediction assuming that all other parameters remain the same as in the reference case. There are two key lessons to draw from the table. First, when the standard deviation of the observation is small enough, most predictions no longer hold for models including learning. The only prediction that is not sensitive to a decrease in the standard deviation of the observation is the prediction for outcome 2 for the pure learning model. Second, if divorce costs are high enough, some predictions no longer hold for the pure changes model and the mixed model.

Since I am interested in comparing empirically observable outcomes of these models, the relevant question is whether we can still distinguish between models when the parameters are outside the range that is compatible with the reference case predictions. Since small values of the standard deviation of the observation are not compatible with standard predictions, I pick a such a small value, i.e. 2, and examine how the predictions differ from the reference case. Parameters other than the standard deviation of the observation remain as in the reference case. Table 5 shows that, under this alternative set of parameters, predictions for outcome 1 are the same for all models. The divorce hazard monotonically decreases with marriage duration even in the pure learning model. This is because, with a small enough standard deviation of the observation, many couples are able to gather sufficient information in one single period to divorce immediately rather than wait to gain more information about marriage quality. The predictions for outcome 3 have also become the same for all models but the mixed model. Only the predictions

for outcome 2 can still allow us to distinguish between the different models. This shows that it is difficult to distinguish between the models if we assume that there is only little learning going on.

While it is difficult to distinguish between the pure learning model and the pure changes model if there is little learning, it is also difficult to distinguish between these two models if divorce costs are high enough. Specifically, if divorce costs are as high as 60 and other parameters stay the same as in the reference case, then the predictions for the pure changes model are the same as for the mixed model. So it is relatively difficult to distinguish between the pure learning model and the pure changes or mixed model when the divorce cost is 60 or higher. However, a divorce cost of 60 is very high, and arguably unrealistic, as it is double the calibrated divorce cost.

Overall, I conclude that it is possible to distinguish empirically between different models under some reasonable parameter restrictions. However, there exist parameter combinations for which the models are difficult to distinguish from each other, and this is especially an issue if there is little learning (low standard deviation of the observation). How much we can learn from the data depends in practice on how the outcomes in Table 2 turn out empirically. I now turn to the empirical analysis.

3 Empirical Evidence: main results

3.1 Data used

I use monthly data from the Survey of Income and Program Participation (SIPP) 1990, 1991, 1992, 1993, 1996, 2001 and 2004¹³. I retain all observations of married persons. The data is reshaped so that the panel identifier corresponds to marriages, not individuals. Thus, for individual characteristics, there is one variable for the woman and one for the man. Marriages for which neither of the partners is ever observed to be the head of the household are not in the sample: these are special cases such as young couples living with

¹³These years are all the years in which a SIPP panel was started.

their parents. Marriages are observed for at most 4 years, and the average window of observation for a marriage is 29 months or a bit more than 2 years. A marriage is defined to end the first time I observe a separation or a divorce. A marriage is right-censored whenever there is any gap in the observations for either partner and that gap is not explained by separation or divorce.

Job loss is defined as either getting laid off or fired. For each category and each spouse, I retain only the first such event observed in the sample. This is to insure better consistency with the model: indeed, the model assumes that marriages have “normal” histories before the shock and so using second shocks would make this assumption less plausible¹⁴. Laid off and fired are two possible responses to the question of why the last job ended; the distinction is therefore based on self-reports.

Table 6 reports summary statistics for the variables of interest. The monthly probability that a marriage ends in separation or divorce is 0.13%, which amounts to 1.56% per year. There are 93,505 marriages in the sample, and of those 3.76% end in divorce or separation during the observation period. Job loss is fairly uncommon: for example, the probability that in any given month the husband has been fired during the past year is 0.27%. Thus, one really does need a large sample to investigate the impact of job loss on divorce, especially since I want to allow for this impact to vary with the duration of the marriage. This is why it makes sense to use the SIPP over alternative datasets such as the NLSY 1979. Also, since job loss in any given month is so rare, it is necessary to define job loss over the past year to get a measure that is not too noisy. Moreover, I hypothesize that, due to transaction costs and delays, it takes at least a few months for a marriage to be dissolved once it has been hit by a fatal negative shock. This justifies looking at job losses that occurred not just in the current month but also within the past year.

¹⁴In practice, more than 90% of observations from marriages with at least one job loss experience only one job loss during the window of observation.

3.2 Econometric specification

Basic specifications relating job loss to the divorce hazard use a Cox proportional hazard model. Such a model assumes that the hazard of separation at marriage duration k is given by:

$$\lambda(k; X) = \lambda_0(k) \exp(\beta' X) \quad (16)$$

where $\lambda_0(k)$ is the baseline hazard estimated non-parametrically, and X is the set of relevant covariates. The Cox model assumes that these covariates have a proportional effect on the baseline hazards.

Since I am also interested in how the impact of job loss varies with marriage duration, I relax the proportionality assumption by allowing the coefficient β to be time-varying. Assume the hazard is given by:

$$\lambda(k; X) = \lambda_0(k) \exp[\beta' + m(k)' X] \quad (17)$$

Assuming that $m(k) = 0$, i.e. that the standard Cox model is correct, then the maximum likelihood estimate of β , $\hat{\beta}$, satisfies:

$$\sum_{i \in D} \left\{ X_i - \frac{\sum_{j \in R_i} X_j \exp(X_j' \hat{\beta})}{\sum_{j \in R_i} \exp(X_j' \hat{\beta})} \right\} = 0 \quad (18)$$

where D is the set of indices for those marriages that end and R_i is the set of marriages that are at risk when marriage i ends. Schoenfeld's (1982) residuals are defined as:

$$\hat{r}_i = X_i - \frac{\sum_{j \in R_i} X_j \exp(X_j' \hat{\beta})}{\sum_{j \in R_i} \exp(X_j' \hat{\beta})} \quad (19)$$

Grambsch and Therneau (1994) showed that $E(\hat{r}_i) \approx V_i g(k_i)$, where V_i is the variance matrix of β , and k_i is the time when marriage i ends. Thus a smoothed plot of $\hat{V}_i^{-1} \hat{r}_i + \hat{\beta}$ versus marriage duration k will reveal the functional form of $\beta(k)$. I am specifically interested in how the impact of job loss varies with marriage duration, so I focus on relaxing

the proportionality assumption for these variables. However, in order to disentangle how the impact of job loss varies with marriage duration, it is important to also allow other control variables to have a time-varying impact. In order to determine if other controls satisfy the proportionality assumption, I test whether the scaled Schoenfeld residual is linearly related to time¹⁵. If some controls violate the proportionality assumption, I stratify on them, or add an interaction with time.

Finally, note that the sample of left-censored marriages is a stock sample with follow-up. Therefore, in my analysis, I take into account the fact these marriages only enter the sample at the date of the first interview¹⁶, and are therefore only at risk of dissolution after that date.

3.3 Results

Figure 5 plots the non-parametric estimate of the divorce hazard. I applied a Kernel-weighted local polynomial smoothing using an Epanechnikov kernel¹⁷. The divorce hazard does not significantly increase starting from 0 duration¹⁸. Looking at the predictions for outcome 1 in Table 2, the observed pattern is most consistent with the pure changes model.

¹⁵I use Stata's `estat phtest` command to do so.

¹⁶I use the `enter` option in Stata's `stset` command to specify this.

¹⁷This was implemented with Stata's `lpoly` command. A similar result is obtained using the Stata command `sts graph, hazard kernel(epan2), degree 1, bandwidth 8`. Both methods are robust to boundary bias (see Cleves et al., 2002, p. 115), and this is important since boundary bias would tend to underestimate the hazard at low marriage durations, thus potentially giving the false impression that the hazard increases with duration. Ultimately, I chose local polynomial smoothing because I want to use the same smoothing for scaled Schoenfeld residuals as for the divorce hazard.

¹⁸This feature of the divorce hazard may surprise the reader since some previous research reported that the divorce hazard initially decreases with duration. To further investigate this issue, I estimated the divorce hazard using the National Longitudinal Study of Youth 1979 (NLSY79) data with interviews up to 2008. The divorce hazard in the NLSY79 is quite similar to the one estimated from the SIPP past the first two years of the relationship. By contrast, in the first two years, the divorce hazard increases when using the NLSY79. I explored the reason for these differences and found that this is not related to differences in demographics between the two samples. Instead, the difference is most likely explained by the fact that the SIPP has quarterly interviews while the NLSY79 interviews participants every year initially, and every two years after 1992. This means that the NLSY79 is likely to miss very short marriages occurring in between interviews. In fact, when using the SIPP data and restricting the sample to marriages that began before the first interview for each individual, I also find that the divorce hazard increases in the first two years. This can again be explained by the fact that using such a stock sample results in an over representation of longer duration marriages.

On the other hand, the confidence intervals around the hazard in Figure 5 are large enough to allow for a small increase in the divorce hazard starting from month 0, so I cannot formally reject models other than the pure changes model. It is only possible to reject that there is a lot of learning going on, i.e. that the standard deviation of the observation is larger than some threshold. Indeed, the sensitivity analysis performed in Table 4 shows that, in the pure learning model, the higher the standard deviation of the observation, the later the maximum of the divorce hazard occurs and the higher the proportional difference between the divorce hazard at period 0 and the maximum of the divorce hazard. Intuitively the divorce hazard has a longer and steeper increasing phase if learning is slower. So I can reject the pure learning model with a large standard deviation of the observation but I cannot reject the pure learning model with a small standard deviation of the observation. Additionally, I cannot reject that the pure learning model, the mixed model and the learning with shocks model make incorrect predictions with respect to the divorce hazard (specifically that the divorce hazard increases initially), but I can reject that the pure changes model makes incorrect predictions with respect to the divorce hazard. Overall, I conclude that the pure changes model is most consistent with the divorce hazard as estimated from the data.

Table 7 estimates the impact of job loss on the separation hazard as a function of how long ago the job loss happened. The first two columns examine the impact of job loss 1 to 12 months ago, while the last two columns estimate the impact of a job loss 13 to 24 months ago. In the first column, without controls, the husband's job loss is associated with a significant and large increase in the divorce hazard. The impact of the husband getting fired is larger than the impact of a lay off but not significantly so. The wife getting fired has a significant and positive impact on the divorce hazard that is comparable in size to the impact of the husband getting fired. By contrast, the wife getting laid off does not significantly affect the divorce hazard. This may be because getting laid off is more closely related to the earnings potential, which matters more for men in their traditional role as breadwinners, whereas getting fired is related to a difficult personality for both

men and women.

Column 2 adds a number of controls reflecting marriage characteristics: age at the beginning of the marriage, race, education, prior marital history. The role of these controls is to allow for the fact that marriages may be governed by different parameters due to heterogeneity in marriage markets. These controls do not include any variable that may be correlated with marriage quality, such as number of kids, home ownership, or pre-marital cohabitation¹⁹. When adding these controls, the coefficients on job loss stay significant but get a little smaller²⁰. Note that, in the column with controls, stratification and interactions with duration²¹ have been used to get rid of non-proportional effects for all right-hand side variables other than job loss (see note to Table 7 for details). The last two columns repeat the specifications for the first two columns, but use job loss 13 to 24 months ago. The impact of job loss 13 to 24 months ago is not significantly different from 0. Thus, using predictions for outcome 2 in Table 2, these results are most consistent with the pure changes model or the mixed model and suggest that match quality changes fairly rapidly. Indeed, using the sensitivity analysis from Table 4, I can show that the higher the standard deviation of the process and the faster the impact of job loss on the divorce hazard converges to 0 as time since job loss increases. Since the estimates of the impact of job loss 13 to 24 months ago have a fairly large standard error, the confidence interval typically includes both positive and negative values, and therefore I cannot reject the predictions from the pure learning model or the learning with shocks model. Still, there is no evidence in favor of these two models, or, more formally, if I assume that the predictions of the models with respect to outcome 2 are false, I cannot reject that they

¹⁹Including such endogenous controls has no effect on the qualitative conclusions from this study

²⁰I also ran a specification controlling for labor market outcomes other than job loss that may be short-run consequences of job loss: monthly earnings in the last year, number of months unemployed during the last year, number of months inactive during the last year. These controls do reduce the magnitude of the coefficients on job loss but job loss conserves its statistical significance.

²¹Using the model without stratifying or interacting any control with a function of marriage duration does not qualitatively affect any of the main results. I stratified on the education variables, and added a time interaction for the two other variables that failed the proportionality test. One of these other variables, the wife's age at the beginning of the marriage, was continuous, therefore not appropriate for stratification. The other variable, prior marriages for the husband, is a dummy that too rarely takes the value of 1 (1.8% of observations) to yield reliable stratified estimates.

are indeed false. Since the learning with shocks model does not find much support, this suggests that job loss does not have a purely causal effect on the divorce hazard. Instead, otherwise worse marriages are more likely to experience job loss²². Overall, I conclude that the pure changes model or the mixed model are most consistent with the evidence from Table 7.

Figure 6 plots the non-parametric estimate of the job loss on divorce as a function of marriage duration. For reference, a linear fit is also plotted. The residuals are from the regression in Table 7, column 2. Using no controls does not affect the shape of the estimates but increases the mean and shrinks confidence bands. These estimates correspond to the log of the hazard ratio and thus allow me to test the predictions for outcome 3 from Table 2. Generally, the impact of job loss increases with marriage duration. The only exception to this general pattern is the wife getting laid off, but then this does not significantly affect the divorce hazard in the first place. Additionally, there is no evidence for a decline in the impact of job loss starting from zero duration. Therefore, the evidence is most consistent with the pure changes model. However, the linear trend is only statistically significant and positive for the husband being fired. For the other job loss events, I cannot reject a slightly negative trend. For the husband being fired, I can therefore reject the prediction of the pure learning model, but not so for other types of job loss. Even for other types of job loss, I can reject that there is a lot of learning, i.e. that the standard deviation of the process is higher than some threshold. Indeed, using the sensitivity analysis from Table 4, I can show that, in the pure learning model, the higher the standard deviation of the observation, the higher the impact of job loss at period 0, and the steeper the decline in the impact of job loss as marriage duration increases from 0. Overall, as for other outcomes, the evidence from Figure 6 is most consistent with the predictions of the pure changes model.

²²This implies that divorce should also predict subsequent job loss. I test this using a sample of jobs from the SIPP and divorce during the previous quarter as a covariate in a Cox proportional hazard model of job separation, which also includes standard covariates such as education. The results (not shown here) show a significant and positive impact of divorce on the probability of a layoff; the impact of divorce on someone getting fired is also positive but insignificant.

More concretely, the reason why the impact of job losses experienced by men tends to increase with marriage duration (Figure 6) is most likely that over time there are more and more marriages whose quality is fairly mediocre (as opposed to really bad), and so this increases the proportion of marriages that are at risk of dissolving if they are hit by a bad shock relative to those marriages that are at risk of dissolving even in the absence of a shock. The psychology literature seems to confirm this hypothesis. Using longitudinal data, Kurdek (2005) shows that marital satisfaction decreases in the first years of marriage, especially for women. It would be important to confirm these findings in the present study. However, the SIPP does not provide any direct information on marital satisfaction. For this reason, I turn to another data set to investigate the plausibility of this hypothesis. The National Survey of Families and Households contains such information. The overall sample for the first wave conducted in 1987-1988 includes a main representative cross-section of 9,637 American households²³. The survey asks married individuals about the date of their marriage, so I can reconstruct their marriage duration as of the date of the interview. The survey also asks questions about marital happiness on a 1-7 scale and the subjective probability of divorce on a 1-5 scale. The latter would seem to best correspond to the theme of this study. However, the subjective probability of divorce has fewer possible values and is much more skewed than marital happiness. For these reasons, I choose to focus on marital happiness. Marital happiness is indeed negatively correlated (-0.48) with the subjective probability of divorce. For example, among women with level 5 of satisfaction, only 50% rate their probability of divorce as very low, compared to 72% at level 6 and 91% at level 7.

In Figure 7, I plot the distribution of women's responses to the question about marital happiness as a function of the duration of their relationship. I find that women who have been married for less than 6 months are most likely to report being "very happy" with their marriage. The proportion of those who declare themselves very happy decreases quite strongly with marriage duration, and this decrease is faster earlier in the relation-

²³The survey also includes an oversampling of blacks, Puerto Ricans, Mexican Americans, single-parent families, families with step-children, cohabiting couples and recently married persons.

ship. A parallel phenomenon is that the proportion of those who declare themselves only “somewhat happy” (category 5) increases quite dramatically with marriage duration. On the other hand, the proportion of those who declare themselves unhappy to any degree does not change over time. This evolution in the distribution of marital bliss can explain why the impact of a husband’s job loss increases with time: in older relationships, wives are less likely to be satisfied and therefore a negative labor market outcome affecting their husband is more likely to dissolve the marriage. Finally, note that the way marital happiness evolves in cross-sections of marriages of different durations is not consistent with what one would expect based on the pure learning model. In the pure learning model, one would expect to see more very unhappy wives at short marriage duration than at long marriage duration. Under the learning model, people wait until they have enough information and only then separate, which means that in the beginning one should see more unhappy wives because they are not yet sure whether they should separate. The pattern in Figure 7 is rather consistent with a model where relationships evolve over time, and tend to slowly regress to the mean (this could be modeled by assuming $\rho < 1$ and $c > 0$ in equation (1)). The evidence in Figure 7 does not preclude that wives are somewhat uncertain about the quality of their relationship, but it is inconsistent with the pure learning model.

Overall, using the predictions from Table 2, I conclude that the pure changes model is most consistent with the data. If I assume that the predictions of the pure changes model are false, I can usually reject that they are false. At the same time, if I assume that the predictions of the pure learning model are false, I can usually not reject that they are false. This holds for the reference case. However, looking at the predictions from Table 5, it is clear that, based on the data, it is not possible to reject that outcomes 1 and 3 are driven by a pure learning model with a small amount of learning. Only outcome 2 from table 5 allows us to distinguish between the pure learning model and the pure changes model, and the pattern in the data for outcome 2 is more favorable to the pure changes model. Additionally, data on the distribution of marital satisfaction at different marriage

durations is not supportive of the pure learning model. Overall, I conclude that, while it is difficult to reject a model with a small amount of learning, the data does not support a large amount of learning (i.e. high standard deviation of the observation). At the same time, patterns in the data are consistent with changes in marital quality driving the bulk of divorces.

4 Conclusion

This paper examines the fundamental reasons underlying the evolution of the divorce probability over the course of a marriage. Although learning about marital quality has been often proposed as an explanation for the divorce hazard, this mechanism finds limited support in the data. On the other hand, the divorce hazard can be fully explained by the assumption that the marital quality follows a random walk. In other terms, divorce can be fully explained by real changes in relationship quality, and without invoking any learning. The fact that learning plays at best a limited role in explaining divorce patterns may be part of the reason why it is so difficult to show empirically that pre-marital cohabitation decreases the probability of divorce. From a policy perspective, the results from this paper suggest that policies that strengthen marriages are those that help couples cope with negative shocks, such as marital counseling or income support policies.

This result is important not only because it sheds light on the substantial mechanisms behind divorce, but also because it clarifies which class of models is most appropriate for marriage. Indeed, learning models are cumbersome, and the mixed model, which also allows for changes in match quality, is even less tractable. If shocks to match quality are the key cause of divorce, then only these shocks need to be modeled in theories whose aim is to explore aspects of the marital relationship other than divorce timing, such as investment in the relationship.

While the theory developed here is suitable to test for the role of learning in marriages, it has important limitations. First, the theory only allows us to test for the presence of

substantial learning. The tests used here do not allow us to reject that there is any learning at all in marriage. Second, I do not explicitly model investments in the relationship. Third, the model does not specify what the sources of shocks to relationship quality are. The empirical work concentrates on one of these shocks, job loss, but other elements also play an important role. Future research should better quantify the relative contribution of various types of shocks to marital dissolution. Another promising research endeavor is to better understand the role of cohabitation. Since learning may not be an important reason for cohabitation, what explains this behavior? Is it for example that partners cohabit instead of marrying because they expect shocks to occur in the near future that could change their valuation of the relationship? Focusing on shocks to relationship quality instead of learning opens interesting avenues for future research on marriage and cohabitation.

References

- [1] ABBRING, JAAP H. AND JEFFREY R. CAMPBELL (2005), “A Firm’s First Year”, Tinbergen Institute Discussion Papers, 05-046/3, Tinbergen Institute.
- [2] ABBRING, JAAP H. AND GERARD J. VAN DEN BERG (2003), “A Simple Procedure for the Evaluation of Treatment Effects on Duration Variables”, IZA Discussion Papers, 810.
- [3] AMATO, PAUL R. (2010) “Research on Divorce: Continuing Trends and New Developments.” *Journal of Marriage and Family* 72 (3) (June 1): 650-666. doi:10.1111/j.1741-3737.2010.00723.x.
- [4] ARULAMPALAM S., MASKELL S., GORDON N., CLAPP T. (2002), “A Tutorial on Particle Filters for Online Nonlinear/Non-Gaussian Bayesian Tracking”, *IEEE Transactions on signal processing*, Vol. 50, No. 2, Feb. 2002.

- [5] AXINN, WILLIAM, AND ARLAND THORNTON (1992), "The Relationship between Cohabitation and Divorce: Selectivity or Causal Influence?" *Demography*, 29(3): 357-74. Reprint No. 386.
- [6] BRIEN, MICHAEL J., LEE A. LILLARD AND STEVEN STERN (2006), "Cohabitation, Marriage, And Divorce In A Model Of Match Quality" *International Economic Review*, vol. 47(2), pages 451-494, 05.
- [7] CASSANDRA A. R. (1998), "Exact and Approximate Algorithms for Partially Observable Markov Decision Processes", Ph.D. thesis, Brown University, May 1998.
- [8] CHARLES K. AND STEPHENS M. "Job Displacement, Disability, and Divorce", *Journal of Labor Economics*, April 2004, vol.22, n.2, p.489-522.
- [9] CLEVES MARIO , WILLIAM W. GOULD, ROBERTO G. GUTIERREZ, AND YULIA MARCHENKO, "An Introduction to Survival Analysis Using Stata, 2nd Edition", Stata Press, 2008.
- [10] CONGER, R.D., RUETER, M.A., ELDER, G.H. JR, *Journal of Personality and Social Psychology*, 1999 Jan, 76(1),54-71.
- [11] FARBER H. S. (1994), "The Analysis of Interfirm Worker Mobility", *Journal of Labor Economics*, Vol 12., 554-593.
- [12] FREDERICK S., LOEWENSTEIN G. AND O'DONOGHUE T. (2002), "Time Discounting and Time Preference: A Critical Review", *Journal of Economic Literature*, Vol 40., No. 2, 351-401.
- [13] GADALLA, TAHANY M. (2008) "Impact of Marital Dissolution on Men's and Women's Incomes: A Longitudinal Study." *Journal of Divorce & Remarriage* 50 (1): 55-65. doi:10.1080/10502550802365714.
- [14] GIBBONS R., KATZ L., "Layoffs and Lemons", *Journal of Labor Economics*, 1991, vol. 9, no. 4, 351-380.

- [15] GRAMBSCH P., THERNEAU T., “Proportional Hazards Tests and Diagnostics Based on Weighted Residuals”, *Biometrika*, Vol. 81, No. 3. (Aug., 1994), pp. 515-526.
- [16] HAUSKRECHT M. (2000), “Value-Function Approximations for Partially Observable Markov Decision Processes”, *Journal of Artificial Intelligence Research*, 13 (2000) 33-94.
- [17] HESS G. (2004), “Marriage and Consumption Insurance: What’s Love Got to Do with it?”, *The Journal of Political Economy*, Vol 112., no.2, 290-318.
- [18] HITSCH G., HORTACSU A., ARIELY D. (2010), “Matching and Sorting in Online Dating”, *American Economic Review*, 100(1), March 2010.
- [19] JOVANOVIĆ B. (1979), “Job Matching and the Theory of Turnover”, *The Journal of Political Economy*, Vol 87., 972-990.
- [20] KAHN L. AND LANGE F. (2010), “Employer Learning, Productivity and the Earnings Distribution: Evidence from Performance Measures”, working paper.
- [21] KARNEY B., BRADBURY T., “Attributions in Marriage: State or Trait? A Growth Curve Analysis”, *Journal of Personality and Social Psychology*, 2000, Vol, 78, No. 2, 295-309.
- [22] KINNUNEN U., FELDT T., “Economic stress and marital adjustment among couples: analyses at the dyadic level”, *European Journal of Social Psychology*, 34, 519-532 (2004).
- [23] KURDEK L., “Gender and Marital Satisfaction Early in Marriage: A Growth Curve Approach”, *Journal of Marriage and Family*, 67 (February 2005): 688-700.
- [24] MARINESCU I. E.(2009), “Job Security Legislation and Job Duration: Evidence from the U.K.”, *Journal of Labor Economics*, vol. 27, no. 3 (p. 475-486).
- [25] MORTENSEN D., PISSARIDES C., “Job Creation and Job Destruction in the Theory of Unemployment” , *Review of Economic Studies* 61 (July 1994) 397-415.

- [26] NAGYPAL E. (2007), “Learning by Doing vs. Learning About Match Quality: Can We Tell Them Apart?,” *Review of Economic Studies*, vol. 74(2), pages 537-566, 04.
- [27] REGE, MARI AND KJETIL TELLE AND MARK VOTRUBA (2007), “Plant Closure and Marital Dissolution”, Discussion Papers, 514, Research Department of Statistics Norway.
- [28] SCHOENFELD D., “Partial residuals for the proportional hazards regression model”, *Biometrika*, 1982, 69, 1, p.239-241.
- [29] SMOCK, PAMELA J., WENDY D. MANNING, AND SANJIV GUPTA. (1999) “The Effect of Marriage and Divorce on Women’s Economic Well-Being.” *American Sociological Review* 64 (6) (December 1): 794-812. doi:10.2307/2657403.
- [30] SULLIVAN D., VON WACHTER T. (2009), “Job Displacement and Mortality: An Analysis Using Administrative Data”, *The Quarterly Journal of Economics*, 124 (3): 1265-1306.
- [31] SVARER M. (2004), “Is Your Love in Vain? Another Look at Premarital Cohabitation and Divorce”, *Journal of Human Resources* 39(2), 523-536.
- [32] WEISS Y., WILLIS R. J. (1997), “Match Quality, New Information, and Marital Dissolution”, *Journal of Labor Economics*, Vol. 15, No. 1, Part 2: Essays in Honor of Yoram Ben-Porath. (Jan., 1997), pp. S293-S329.

Figure 1: Timing of partner's decisions

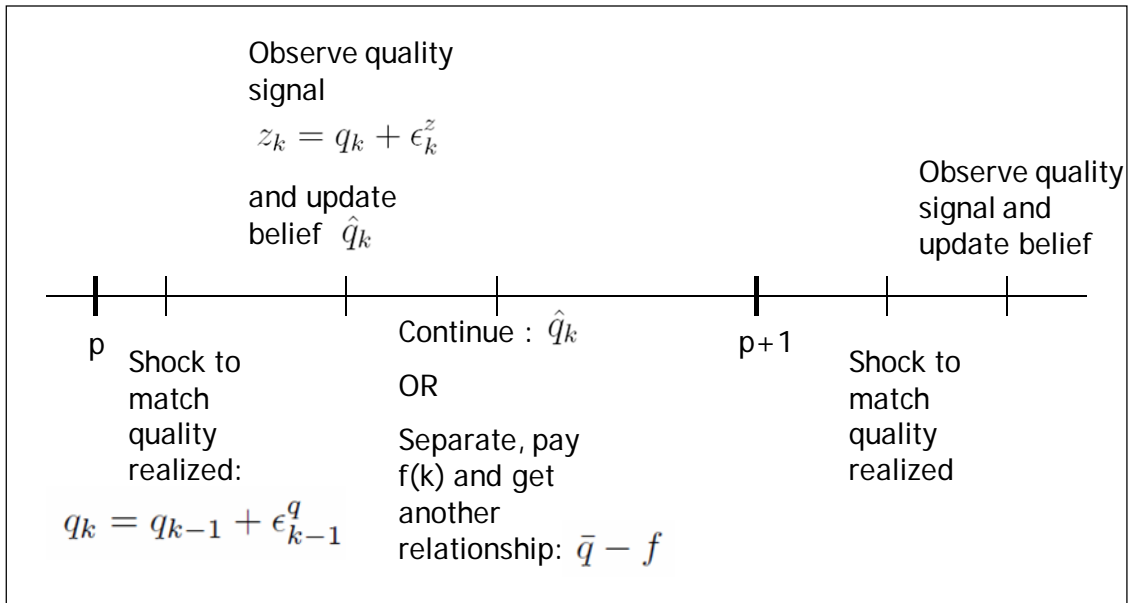



Table 1: Reference parameters

	Pure learning model	Pure changes model	Mixed model	Learning with shocks
<u>Parameters of interest</u>				
Standard deviation of observation	10	0	10	10
Standard deviation of process	0	5	5	0
Probability of job loss occurring	N/A	N/A	N/A	0.05
Probability of job loss state ending	N/A	N/A	N/A	0.1
Standard deviation of noise to the observer*	N/A	5*	N/A	N/A
<u>Parameters held constant</u>				
Mean of prior	30	30	30	30
Standard deviation of prior	5	5	5	5
Drift of process	0	0	0	0
Auto-correlation of process	1	1	1	1
Threshold for bad observation	15	15	15	15
Separation cost	30	30	30	30
Discount factor	0.8	0.8	0.8	0.8
<u>Technical parameters</u>				
Range of match qualities	[0,60]	[0,60]	[0,60]	[0,60]
Number of match quality values	801	801	801	801
Maximal duration	50	50	50	50


Note: the standard deviation of noise to the observer is only used in Figure 4.

Figure 2: Theoretical divorce hazard under alternative models




Note: The parameters for the calculation of these hazards can be found in Table 1.

Figure 3: Theoretical log hazard ratios under alternative models



Note: The parameters for the calculation of these hazards can be found in Table 1. The hazard ratio is defined as the hazard if job loss occurred at k divided by the hazard if no job loss occurred at k . Time k is represented on the x axis. The log ratios for the pure learning model and for learning with shocks have been smoothed with a moving average with a span of 5 for all values of k greater or equal to 10 in order to attenuate discretization artifacts.

Figure 4: Theoretical log hazard ratios after job loss occurred at period 5



Note: The parameters for the calculation of these hazards can be found in Table 1. The hazard ratio is defined as the hazard at k if job loss occurred at period 5 divided by the hazard at k if no job loss occurred at period 5. Time k is represented on the x axis.

Table 2: Empirically testable predictions using reference parameters

Outcome Number	Outcome description	Prediction under alternative models			
		Pure learning	Pure changes	Mixed model	Learning with shocks
1	Slope of the divorce hazard (see Fig. 2)	+ then -	-	+ then -	+ then -
2	Sign of the impact of job loss on the divorce hazard a few periods after job loss occurred (see Fig. 4)	+	Undefined in simple model. In the model with extra noise, + then approaches 0 when the number of periods since job loss is large enough	+ then approaches 0 when the number of periods since job loss is large enough	+, then - when the number of periods since job loss is large enough
3	Slope of the impact of job loss as a function of marriage duration (see Fig. 3)	-	+	- then +	- then +

Note: the reference parameters are in Table 1.

Table 3: Range of parameters tested

Standard deviation of observation	Standard deviation of process	Separation cost	Threshold for bad observation
0	0	20	10
1	1	30	15
2	2	40	20
3	3	60	25
4	4		
5	5		
10	7		
20	10		

Note: I have computed the outcomes in Table 2 under all possible combinations of the parameters above. Parameters not listed above remain the same as in the reference case.

Table 4: Sensitivity of predictions to changes in parameters

		Range of parameters yielding the same predictions as in the reference case			
Model	Outcome	S.d. of process	S.d. of observation	Divorce cost	Threshold for bad observation
Pure learning	1	0	[4:20]	[20:60]	N/A
	2	0	[1:20]	[20:60]	[10:25]
	3	0	[3:20]	[20:60]	[10:20]
Pure changes	1	[1:10]	0	[20:30]	N/A
	2	[1:10]	0	[20:60]	[10:25]
	3	[1:10]	0	[20:30]	[10:15]
Mixed	1	[1:10]	[3:20]	[20:60]	N/A
	2	[1:10]	[3:20]	[20:60]	[10:25]
	3	[2:10]	[2:20]	[20:40]	[10:20]
Learning with shocks	1	0	[4:20]	[20:60]	N/A
	2	0	[3:20]	[20:60]	N/A
	3	0	[4:20]	[20:60]	N/A

Note: For each parameter, the range given is the range that preserves the reference case predictions in Table 2, assuming that all other parameters remain the same as in the reference case. Only parameter values in Table 3 were tested.

Table 5: Predictions under a low level of learning

Outcome Number	Outcome description	Prediction under alternative models			
		Pure learning	Pure changes	Mixed model	Learning with shocks
1	Slope of the divorce hazard	-	-	-	-
2	Sign of the impact of job loss on the divorce hazard a few periods after job loss occurred	+	Undefined in simple model. In the model with extra noise, + then approaches 0 when the number of periods since job loss is large enough	Undefined in simple model. In the model with extra noise, + then approaches 0 when the number of periods since job loss is large enough	-
3	Slope of the impact of job loss as a function of marriage duration	+	+	- then +	+

Note: Parameters as in the reference case (Table 1), except that the standard deviation of the observation is set to 2.

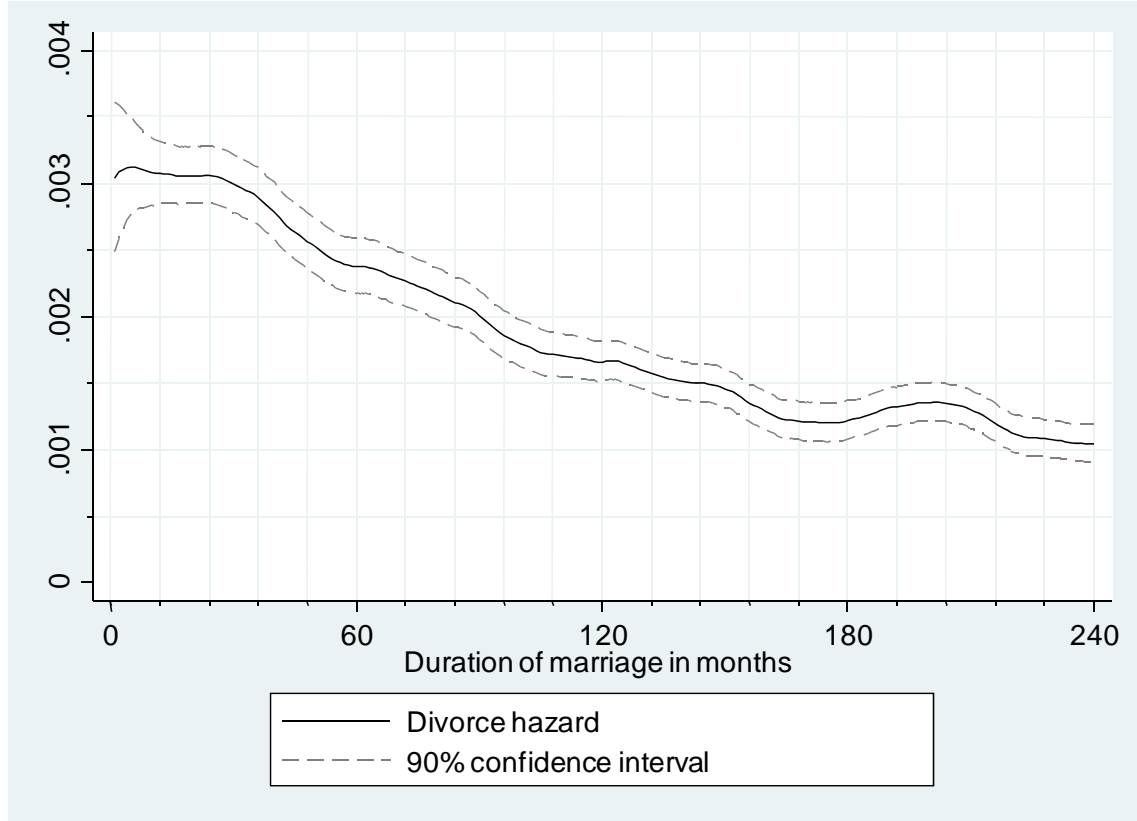
Table 6: Summary statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
Divorce or separation	2797181	0.001256	0.035417	0	1
Husband laid off in the last year	2797181	0.011647	0.107291	0	1
Husband fired in the last year	2797181	0.00267	0.051605	0	1
Wife laid off in the last year	2797181	0.008068	0.08946	0	1
Wife fired in the last year	2797181	0.001847	0.04294	0	1
Husband has had another marriage	2797181	0.018288	0.133991	0	1
Wife has had another marriage	2797181	0.017493	0.131098	0	1
Age of the husband at the beginning of marriage	2796877	28.47455	9.432294	12	87
Age of the wife at the beginning of marriage	2796705	25.95334	8.744819	12	87
Husband is 5 years older than the wife or more	2797181	0.191416	0.393416	0	1
Wife is 5 years older or more	2797181	0.03128	0.174073	0	1
Husband is white	2797181	0.888015	0.315348	0	1
Both same race	2797181	0.973222	0.161433	0	1
Both high school educated or less	2797181	0.374762	0.484062	0	1
One with high school and one with some college or more	2797181	0.239599	0.426839	0	1
Both some college or more	2797181	0.263985	0.440791	0	1

Note: An observation is a marriage*month.

Source: Survey of Income and Program Participation, 1990, 1991, 1992, 1993, 1996, 2001, 2004.

Figure 5: Divorce hazard



Notes: Kernel-weighted local polynomial smoothing using an Epanechnikov kernel, degree 1, bandwidth 8. The smoothing was achieved using durations up to 300 but the hazard is only graphed up to duration 240.

Source: Survey of Income and Program Participation, 1990, 1991, 1992, 1993, 1996, 2001, 2004.

Table 7: Impact of job loss on divorce as a function of time elapsed since the job loss

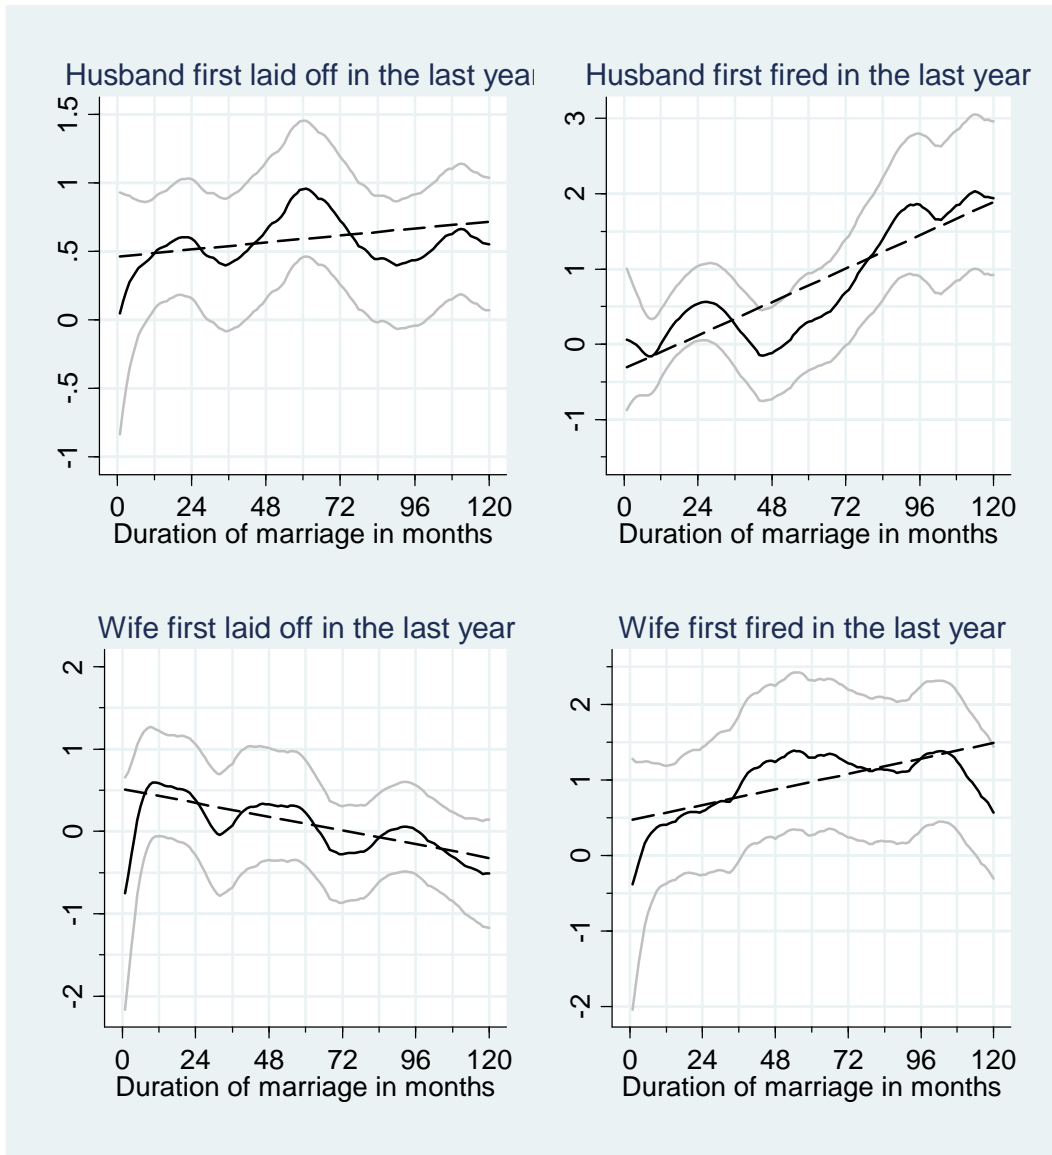
	Job loss 1-12 months ago		Job loss 13-24 months ago	
	No controls	Controls	No controls	Controls
Husband				
Laid off	0.730*** (0.126)	0.572*** (0.126)	0.194 (0.232)	0.0428 (0.232)
Fired	0.887*** (0.206)	0.652*** (0.206)	0.421 (0.380)	0.212 (0.380)
Wife				
Laid off	0.188 (0.194)	0.146 (0.194)	0.404 (0.252)	0.335 (0.252)
Fired	1.046*** (0.225)	0.902*** (0.225)	0.445 (0.449)	0.233 (0.449)
Observations	820,845	820,799	421,453	421,428

Standard errors in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%

Notes: Cox proportional hazard model for the marriage ending in a separation or divorce. Columns 2 and 4 use a stratified Cox model. The controls that are stratified on are: a dummy taking the value 1 if only one of the partners is a high school dropout or if both partners are high school graduates, a dummy taking the value 1 if one of the partners is a high school graduate and the other has some college education, and a dummy taking the value 1 if both partners have some college education. The other controls in columns 2 and 4 are: a dummy for having contracted a previous marriage (one dummy for each partner), a dummy for the husband having contracted a previous marriage interacted with marriage duration, age at the beginning of the marriage (one variable for each partner), wife's age at the beginning of the marriage interacted with marriage duration, age difference between the partners (set of 2 dummies: wife older by 5 years or more, husband older by 5 years or more), white husband dummy, partners of the same race dummy.

Source: Survey of Income and Program Participation, 1990, 1991, 1992, 1993, 1996, 2001, 2004.


Figure 6: Linear and non-parametric estimates of the impact of job loss as a function of marriage duration



Notes: Cox proportional hazard model for the marriage ending in a separation or divorce. The graph plots a local polynomial smooth (degree 1, bandwidth 8, Epanechnikov kernel) of scaled Schoenfeld residuals for various definitions of job loss together with 90% confidence intervals in gray; the dashed line is a linear fit of the residuals on time. The specification used is the same as in Table 7, column 2, but includes marriages up to 15 years old.

Source: Survey of Income and Program Participation, 1990, 1991, 1992, 1993, 1996, 2001, 2004.

Figure 7: Distribution of marital happiness for women, at varying relationship durations



Notes: 1 corresponds to “very unhappy” and 7 to “very happy”.
Source: National Study of Families and Households 1987-1988, cross-section.