A reinforced urn process modeling of recovery rates and recovery times

Cheng, Dan; Cirillo, Pasquale

DOI
10.1016/j.jbankfin.2018.08.014

Publication date
2018

Document Version
Final published version

Published in
Journal of Banking & Finance

Journal of Banking and Finance 96 (2018) 1–17

Contents lists available at ScienceDirect

Journal of Banking and Finance

journal homepage: www.elsevier.com/locate/jbf

A reinforced urn process modeling of recovery rates and recovery times

Dan Cheng, Pasquale Cirillo*

Applied Probability Group, Delft University of Technology (TU Delft), Van Mourik Broekmanweg 6, Delft 2628XE, The Netherlands

ARTICLE INFO

Article history:
Received 12 July 2017
Accepted 20 August 2018
Available online 29 August 2018

Keywords:
Loss-given-default
Recovery rate
Reinforced Urn Process (RUP)
Survival analysis
Machine learning
Bayesian methods

ABSTRACT

Answering a major demand in modern credit risk management, we propose a nonparametric survival approach for the modeling of the recovery rate and the recovery time of a defaulted counterparty, by introducing what we call the Recovery Reinforced Urn Process, a special type of combinatorial stochastic process.

The new model allows for the elicitation and exploitation of prior knowledge and experts' judgements, and for the constant update of this information over time, as soon as new data become available. We show how to use it to perform Bayesian nonparametric prediction about the recovered amounts and the (total) recovery time of a series of defaulted exposures.

An application to real data is provided using the Single Family Loan-Level Dataset by Freddie Mac.

1. Recovery modeling: introduction and main issues

In the wake of the latest financial regulations for the banking sector (BCBS, 2011, IFRS9, 2014, EBA, 2017), we provide an answer to the increasing demand of models for the recovery process of defaulted exposures. Differently from most contributions in the literature that mainly focus on the recovered amounts, here we propose a way to also study the duration of the recovery process, a common "known unknown" (Jorion, 2009) in risk management, with the final goal of predicting the possible recovery trajectory of a counterparty, not only on the basis of the available data, but also with the possibility of using experts' judgements and other a priori knowledge to mitigate historical bias (Knight, 1964, Taleb, 2007), at least partially (Shackle, 1955, Derbyshire, 2017). The model we propose is able to learn, improving its performances over time, using the mechanism of Bayesian update, or machine learning in computer science language.

When a counterparty defaults, the corresponding loss is not necessarily equal to the Exposure-at-Default (EAD), that is the nominal value of the exposure at the time of default. In fact, thanks to the recovery process, i.e. the set of all the procedures that can be put in place to collect the amount due, one may be able to recover at least a part of the outstanding exposure (Hull, 2015). But these procedures are costly and may require a substantial amount of time, a very important variable for correct risk management.

The recovery rate (RR) is the amount of principal and accrued interest on a defaulted exposure that can be recovered, expressed as a percentage of its face value, once again the EAD. In what follows, to ease notation, we will consider $RR \in [0\%, 100\%]$, even if the case $RR > 100\%$ is actually possible, thanks to fees and interests, as frequent among leasing contracts (Schmit, 2004, Zhang and Thomas, 2012). The recovery rate is naturally linked to the Loss Given Default, or *LGD*, that is the percent loss experienced when a counterparty defaults and no further recovery is possible. For $RR \in [0\%, 100\%]$, we have $LGD = 100 - RR$. When $RR \geq 100\%$, we set $LGD = 0\%$.

In the last decade, because of the Basel Accords (BCBS, 2005, 2006, 2011), and the new International Financial Reporting Standards (IFRS9, 2014), recovery modeling has become a major concern for banks and regulators. In particular, the new IFRS 9 regulations–just entered in full force–define LGD assessment and back-testing as a major task for banks and financial institutions (Reitgruber, 2015). Recovery risk, defined as the risk associated with the recovery process, in terms of time length and quantification of the actual loss, is officially one of the new challenges in credit risk management (Schuermann, 2004, Bade et al., 2011, Reitgruber, 2015).

The recovery process is not at all a simple object (Altman et al., 2005a), given that the entire process is influenced by a series of important factors, which affect its success and duration. As Resti and Sironi (2007) point out, an effective recovery depends

* Corresponding author.
E-mail addresses: D.Cheng@tudelft.nl (D. Cheng), P.Cirillo@tudelft.nl (P. Cirillo).

https://doi.org/10.1016/j.jbankfin.2018.08.014
0378-4266/

on the characteristics of the exposure (presence of collateral, degree of effectiveness), of the counterparty (e.g. industry, country and legal framework), on macroeconomic factors like the state of the economy, and on internal factors like the efficiency of a bank in recovering its money, for instance in dealing with out-of-Court settlements.

The complexity of the recovery process is evident if we take into account its possible outcomes, and the fact that its actual duration, the recovery time T, is not known until the very end. We can clearly distinguish three main scenarios: 1) the past due receivable is fully collected, so that, at some random time T, the recovery process terminates with $RR = 100\%$; 2) the uncollected receivable is fully or partially written-down or written-off, because it has no residual market value nor it can be monetized, so that at time T we have $RR < 100\%$; 3) the recovery time exceeds a given maximum time $T^{\max}$, legally imposed, therefore $RR < 100\%$ and $T = T^{\max}$. From a statistical point of view, this last case involves censored data (Kleinbaum and Klein, 2012), something that needs to be taken into consideration (Resti and Sironi, 2007).

The quantification of the recovery rate is a further problem, because also the recovered amount is precisely known only at the end of the recovery process. And while it is true that for the bond issues of large corporations one can typically rely on the so-called market recoveries (computed as the ratio of the price at which the defaulted bond is traded some days after default, provided the market is not illiquid, to the price of that asset at the time of default), it is also true that for smaller counterparties, or for other products like loans, the information about prices is either not available nor reliable. Khieu et al. (2012) have shown that post-default prices tend not to be rational, but rather biased and inefficient. Even when prices are available, the discounted value of cash flows is the preferred measures by both analysts (Frye, 2005) and regulators (BCBS, 2011, IFRS9, 2014).

The paper is organized as follows. Section 2 provides a brief overview of the state-of-the-art in recovery modeling. Section 3 gives an intuitive representation of the recovery process, which is propaedeutical to the model we introduce in Section 4. In Section 5 we provide an extensive application of the model to real data, showing how to deal with the recovery trajectories of defaulted loans in the Single Family Loan-Level Dataset by Freddie Mac (2017). Section 6 closes the paper. All the mathematical contents related to the new model, from theorems to proofs and simulation details, are collected in the Appendix.

2. The state-of-the-art of recovery modeling

The literature on recovery modeling is not as extensive as the one related to the probability of default (PD), but especially for the recovery rates it is nevertheless rich and varied, from the pioneering works on hazard functions for loss data (Kiefer, 1988), to the very recent contributions about the methodological implications of the new regulations on the estimation of LGD (Reitgruber, 2015).

For good reviews, we refer to Altman et al. (2005a), Höochstötter and Nazemi (2013), Gürtler and Hibbeln (2013) and Yashkir and Yashkir (2013). In particular, the first paper also contains details about some less used methods like those based on utility theory.

Statistically, the modeling of LGD includes parametric and nonparametric approaches (Hartmann-Wendels et al., 2014). In the first class, we find the different flavors of generalized linear models, from the convenient beta regression of Huang and Oosterlee (2011) to the logit, probit and tobit regressions described in Bellotti and Crook (2012), or in Loterman et al. (2012), where interesting comparisons are discussed. In a very recent paper, Krüger and Rösch (2017) deal with quantile regression, and they are able to deal with both bulk and tail events, that is with the entire LGD spectrum. Regarding nonparametric methods, NP regression and model trees (Hartmann-Wendels et al., 2014) are common tools, together with mixtures (Calabrese and Zenga, 2010) and beta kernels (Renault and Scaillet, 2004).

Recently, a series of survival analysis models have also been proposed, e.g. Im et al. (2012), according to a trend in PD modeling, but also following the new regulations (BCBS, 2011, IFRS9, 2014) asking for both point-in-time (PIT) and through-the-cycle (TTC) estimates (Chawla et al., 2015, Reitgruber, 2015). We can cite Bonini and Caivano (2013) and Witzany et al. (2012) as examples of pure survival approaches. In particular, in the latter, different models are considered, including interesting semi-parametric and pseudo-survival constructions. For a comparison between regression and survival methods, we refer to Zhang and Thomas (2012).

Most of the contributions in the literature, a notable exception being Bastos (2010), do not present models that are able to learn from data, updating their parameters, and improving their performances over time, whenever new pieces of information become available. This is actually our ambition: using a specially conceived reinforced urn process (Muliere et al., 2000), we propose a nonparametric survival approach that allows for the elicitation and exploitation of prior information, and its automatic updating and correction over time using data. A model that can be seen as a mixture of combinatorial stochastic processes and machine learning, but that, differently from traditional machine learning (Murphy, 2012), gives us full control of its probabilistic features.

3. Visualizing recovery

Without loss of generality, we can discretize the recovery rate of a counterparty in terms of a scale of recovery levels from 0 to m. Level 0 corresponds to no recovery ($RR = 0\%$), levels 1 to $m-2$ represent intermediate stages of recovery ($0\% < RR < 100\%$), while level $m-1$ is full recovery ($RR = 100\%$). Level m is a special stage, not associated to any specific recovery rate, but representing the termination of the recovery process. We can read it as the situation in which the recovery process has reached full recovery or a write-off, and we need some little extra time for closing the bureaucratic procedures. As we shall see, the termination level m is a little artifice needed to fully develop the model, guaranteeing the property of recurrence (Appendix: Lemma 1).

Let us consider a simple example. Set $m = 4$ and define 4 possible levels for the recovery rate:

$$0 : 0\%,\ 1 : (0\%, 50\%],\ 2 : (50\%, 100\%),\ 3 : 100\%, \tag{1}$$

plus the extra level 4 representing the termination of the recovery process.

The larger m, the finer the partition for the discretized recovery rate. Notice that it is not required that each recovery level guarantees the same additional recovery. In other words, in the previous example, we could define a scale in which level 0 corresponds to 0%, level 1 to (0%, 25%], level 2 to (25%, 100%), and level 3 to 100%, according to our needs. The only constraints are that level 0 is equal to 0%, no recovery, level $m-1$ corresponds to 100%, full recovery, and level m is the termination level.

Following intuition, we can take the recovery level to be a nondecreasing quantity: once we have recovered 20% of the outstanding exposure, we can further increase our recovery level or not, but we cannot go back to 10%.

We set time t to be a nonnegative integer ($t \in \mathbb{N}_0$), representing a particular time unit, like days or months. For the recovery process, discrete time is not a limitation: according to Altman et al. (2005a), it is actually more realistic than a continuous time framework.

With the couple (t, l), we can therefore indicate that a given counterparty at time t is at recovery level $l \in \{0, 1, \dots, m\}$.

By definition, the recovery process of a counterparty starts in (0,0), at default, where the recovery clock is yet to start and the recovered amount is null. From (0,0), time starts running and we can spend several time units at a given recovery level, before being able to reach a higher one.

Let us consider a counterparty A, whose recovery trajectory R_A looks like

$$R_A = \{(0,0),(1,0),(2,0),(0,1),(1,1),(0,3),(1,3),(0,4),(1,4)\}. \quad (2)$$

This means that, after default in (0,0), counterparty A spends 2 time units in level 0, where no amount is recovered, visiting (1,0) and (2,0). It then jumps to recovery level 1, where it stays 1 time unit, recovering up to 50% of its EAD in the scale of Eq. (1). Finally, for one time unit, A reaches level 3, full recovery, before jumping to termination level (0,4), which closes the recovery process. For the termination level, in what follows, we always assume a fixed fictitious permanence of 1 time unit (more later). The total recovery time for counterparty A is therefore $T_A = 2 + 1 + 1 = 4$. The time spent in level 4, the termination level, is not taken into consideration for the computation of the total recovery time.

4. Recovery modeling and reinforced urn processes

Consider a portfolio of k defaulted exposures, which we assume exchangeable.[1] Exchangeability is a relaxation of the stronger assumption of independence: the order in which we observe our counterparties is for us irrelevant, and the joint distribution of their recovery times and levels is immune to permutations in the order of appearance. Exchangeability is a common assumption in credit risk modeling, for instance in the class of mixture models (McNeil et al., 2015), where it probabilistically represents the idea that we can split a group of counterparties into subgroups that are homogeneous in terms of risk, that is exchangeable within.

In this paper, the exchangeability of counterparties consists in the exchangeability of their recovery trajectories: it is not important if the recovery process of counterparty i is observed before or after that of counterparty j, because the joint distribution of their recoveries is unaffected, and so our estimates. We assume that the recovery trajectory R_i of counterparty i is representable as in the example of Eq. (2), for $i = 1, \dots, k$.

4.1. The general RUP construction

To develop our model we make use of Reinforced Urn Processes (RUPs) that are a class of combinatorial stochastic processes introduced by Muliere et al. (2000).

A RUP $\{X_n\}_{n\geq 0}$ is characterized by the following elements:

- A countable state space S of all the possible states the process $\{X_n\}$ can visit with positive probability.
- Every element $s \in S$ is associated with a Pólya urn, i.e. an urn characterized by sampling with reinforcement (Mahmoud, 2008), containing colors belonging to a set C, with cardinality $\# C > 0$.
- For each urn we define a function U: $s \in S \to \mathbb{R}^{\#C}$ describing the composition of the urn itself. In other words, for every $s \in S$, urn $U(s)$ contains $N_s(c) \geq 0$ balls of color $c \in C$. To avoid degenerate cases, we assume $\Sigma_{c\in C} N_s(c) > 0$, so that every urn contains a positive amount of at least one color.[2]
- A law of motion $q: S \times C \to S$ indicating how the sampling of the different urns drives the process $\{X_n\}$. Given a color $c \in C$ and two states s, $w \in S$, so that w can be visited from s with positive probability after sampling c from $U(s)$, the function q is such that $w = q(s, c)$. Clearly the definition of q allows for the construction of very different processes, all falling within the general RUP framework of Muliere et al. (2000) to which we refer for details.

Without any loss of generality, fix an initial state s_0. A RUP $\{X_n\}$ on S with initial state s_0 is defined recursively as follows: set $X_0 = s_0$, and for all $n \geq 1$, if $X_{n-1} = s_{n-1} \in S$, sample a ball from the urn $U(s_{n-1})$ associated with s_{n-1}. Now, register the color of the ball, say $c \in C$, put it back in $U(s_{n-1})$, and Pólya-reinforce the urn with $r > 0$ balls of the same color. This increases the probability of picking again that color in the future, if the urn is sampled again: the higher r the stronger the update. Finally, using the rule of motion q, set $X_n = q(s_{n-1}, c)$. The sequence $\{X_n\}$ is a reinforced urn process with initial state s_0 and reinforcement r.

A RUP is just a reinforced random walk on a state space of urns. It is a Bayesian nonparametric model, in which the initial composition of the urns defines the a priori (the way in which the a priori is elicited is given in Appendix), which we can update over time thanks to the sampling of the urns, using reinforcement and other technical conditions we shall discuss later. This possibility of embedding prior knowledge and to learn from data is one of the points of strength of RUPs that, in the last years, have been used to develop some interesting models in finance (Cirillo et al., 2010, Peluso et al., 2015), biostatistics (Mezzetti et al., 2007) and other fields.

4.2. The recovery RUP (R-RUP)

A generic RUP can be adapted to model recovery rates and recovery times, we call this special process the Recovery RUP or R-RUP.

We better specify the space S by setting $S = \mathbb{N}_0 \times L$, so that it contains all the couples (t, l) of recovery times t and recovery levels l the process $\{X_n\}$ can visit. As per Section 3, levels 0 to $m-1$ are a discretization of the recovery rates from 0% to 100%, while level m represents the termination level.

The set of colors is now $C = \{c_0, c_1, \dots, c_m\}$, where each color from c_0 to c_m corresponds to a given recovery level, as if we color them.

For the fundamental rule of motion q, we define three behaviors:

A. $q((t,l), c_i) = (t+1, l)$, for $i \leq l \leq m$. In words: if, from the urn centered in (t, l), we extract a ball whose color corresponds to level l or lower, the process moves to the next time unit $t+1$, but stays at the same recovery level l. From (t, l) we move to $(t+1, l)$. The rule can be further simplified if we assume that urns at level l only contain balls of colors $(c_l, c_{l+1}, \dots, c_m)$, so that we can set $q((t,l), c_l) = (t+1, l)$. This is actually the version we adopt from now on.

B. $q((t,l), c_i) = (0, i)$, for $l < i \leq m$. The process jumps to level $i > l$ while time is reset to 0. Resetting time at every jump helps in counting the time units the R-RUP spends in each recovery level it visits. When a c_m ball is extracted from the urn in state (t, l), the recovery process jumps to the termination level. Clearly, if c_m is extracted in level $m-1$, the recovery process has reached full recovery, otherwise, for $l = 0, \dots, m-2$, we are experiencing a write-off.

[1] Given two random variables X_1 and X_2, they are exchangeable whenever $P(X_1 \leq x_1, X_2 \leq x_2) = P(X_1 \leq x_2, X_2 \leq x_1)$, that is when their joint distribution is immune to permutations in the order of the variables. The concept is easily extended to higher dimensions (de Finetti, 1975).

[2] Notice that $N_s(c)$ is a real number, so that we can have 1.3 balls of color c. A real number of balls can be represented using balls of different radius: the bigger the ball, the easier to sample it. Naturally we can choose $N_s(c)$ to be an integer: this will not affect the model, but it can help to have a simpler intuition of the sampling scheme.

C. $q((1,m),c_m)=(0,0)$. The process $\{X_n\}$ can stay in level m only for one time unit thanks to rule A[3], after which it reaches (1, m) and it is reset to (0,0). This last movement allows us, using Lemma 1 in the Appendix, to define a recurrent R-RUP.

To avoid degenerate situations and to facilitate the application to real data, we can impose extra restrictions on the urn compositions. From one side, we can set $N_{(0,l)}(c_i)=0$ for $i=l+1,\dots,m$, and all $l\in L$, guaranteeing that we cannot visit two recovery levels at the same (discrete) time, because when the process touches level l in (0, l), it must spend at least one time unit at this level. From the other side, for all $t\geq 1$, we can impose $N_{(t,m-1)}(c_i)=0$ for $i<m$, and $N_{(t,m-1)}(c_m)>0$, so that the process necessarily moves to the termination level after visiting full recovery for one time unit.

In the R-RUP construction, the recovery process of counterparty i, with all its intermediate stages, can be represented as a sequence of points (t, l) in the space S. For a first counterparty we might for example observe

$$\begin{aligned} R_1 = \quad & [X_0, X_1, \dots, X_{n_1}] = \\ & [(0,0),(1,0),\dots,(t_0,0),(0,l_1),(1,l_1),\dots,(t_1,l_1), \\ & \dots,(0,l_{\max}),\dots,(t_{\max},l_{\max}),(0,m),(1,m)], \end{aligned}$$

where $l_1,l_2,\dots,l_{\max}$ are the levels visited by the counterparty ($l_{\max}<m$ being the maximum recovery level reached), and $t_1,t_2,\dots,t_{\max}$ are the corresponding sojourn times, whose sum represents the total recovery time T_1. The quantities $l_{\max}$ and T_1 fully summarize the recovery trajectory (and risk) of exposure 1.

Similarly to R_1, the recovery process of each counterparty i is thus represented as a block R_i of visited states starting with (0,0). If we consider k exposures, we will have k blocks. Every time we observe (0,0) in the sequence generated by $\{X_n\}$, it means that we are looking at a new counterparty. Since we assume the recovery trajectories of the exposures to be exchangeable, the blocks R_i are exchangeable. The R-RUP is therefore able to model any number of counterparties if, every time a recovery process is over, we reset $\{X_n\}$ to (0,0), so that we can start a new block. In probabilistic terms, $\{X_n\}$ needs to be a recurrent, as we discuss in the Appendix.

In general, after n time steps, for $\{X_n\}$ we have

$$\{X_0,X_1,X_2,\dots,X_n\}=\{R_1,R_2,\dots,R_k\},$$

so that the first n states visited by $\{X_n\}$ can be collected into k 0-blocks, where the number k depends on the number of (0,0) in the sequence of states generated by $\{X_n\}$. A group of k defaulted counterparties will thus be represented by k 0-blocks: $R_1,\dots,R_k$. Naturally the last recovery trajectory R_k could be incomplete at step n, depending on n and k.

4.3. A clarifying example

The maximum recovery level is $m=4$, so that each urn in S is characterized by a set of 5 colors $C=(c_0,c_1,c_2,c_3,c_4)$. According to Subsection 4.2, every urn at recovery level $l\in\{0, 1, 2, 3, 4\}$ only contains balls of color $(c_l,\dots,c_4)$.

We start with the default of exposure 1, which begins its recovery process in $X_0=(0,0)$. We sample urn $U((0, 0))$ and extract a c_0 ball, so that we move to $X_1=q((0,0),c_0)=(1,0)$, that is exposure 1 stays at the recovery level 0 for one time unit. Then, we sample urn $U((1, 0))$ and we extract c_2, therefore we jump to $X_2=(0,2)$, i.e. after staying at recovery level 0 for 1 time unit, the counterparty jumps immediately to level 2, without touching level 1. In urn $U((0, 2))$, we sample again a c_2 ball, we set $X_3=(1,2)$, and we stay at recovery level 2 another time unit. Finally, we extract a c_3 ball from urn $U((1, 2))$ and the exposure reaches the full recovery level 3. The recovery process goes back to state (0,0) after we stay one more time in the full recovery level 3, and in the termination level 4 for "bureaucratic" reasons. Summarizing, for exposure 1 we observe:

$$\begin{aligned} R_1 &= \{X_0,\dots,X_7\} \\ &= \{(0,0),(1,0),(0,2),(1,2),(0,3),(1,3),(0,4),(1,4)\}. \quad (3) \end{aligned}$$

The total recovery time for exposure 1 is $T_1=1+1+1=3$, given that the one-period permanence in level 4 is not counted. The recovery trajectory is visible in Fig. 1(a).

All the urns in S are Pólya, therefore every ball sampled from an urn is reinforced with r extra balls of the same color, thus changing the composition of that urn. Hence, after R_1 has been observed, the probability that a new counterparty follows the same trajectory increases. The R-RUP thus remembers and learns from what happens.

Let us now continue our sampling from urn $U((0, 0))$, where the process $\{X_n\}$ has landed after the resetting due to the rule of motion q applied in $(1, 4)\in R_1$. Imagine that we obtain two further recovery trajectories for exposures 2 and 3, i.e.

$$\begin{aligned} R_2 &= \{X_8,\dots,X_{17}\} \\ &= \{(0,0),(1,0),(2,0),(3,0),(0,1),(1,1),(0,2),(1,2), \\ &\qquad (0,4),(1,4)\}, \end{aligned}$$

and

$$\begin{aligned} R_3 &= \{X_{18},\dots,X_{33}\} \\ &= \{(0,0),(1,0),(2,0),(3,0),(4,0),(0,1),(1,1),(0,2), \\ &\qquad (1,2),(2,2),(3,2),(4,2),(5,2),(6,2),(0,4),(1,4)\}. \end{aligned}$$

Notice that for exposure 2 (and also for 3), because of the recovery trajectory of exposure 1 and the Pólya-reinforcement mechanism, the probability of picking a c_2 ball in (1,0) is higher than what originally experienced by exposure 1 (there are r extra c_2 balls now), and this is true for all the already visited states.

The recovery processes for exposures 1, 2, and 3 can thus be represented as follows:

$$\{R_1,R_2,R_3\}=\{X_0,X_1,\dots,X_{33}\}=\left\{\begin{aligned} &\overbrace{(0,0),\dots,(1,3),(0,4),(1,4)}^{\text{exposure 1}}, \\ &\times\overbrace{(0,0),\dots,(1,2),(0,4),(1,4)}^{\text{exposure 2}}, \\ &\times\overbrace{(0,0),\dots,(6,2),(0,4),(1,4)}^{\text{exposure 3}} \end{aligned}\right\}.$$

Using reinforcement, the R-RUP modifies its transition probabilities at every cycle. In the R-RUP, the initial compositions of the urns represent our a priori, which we modify through sampling, in order to get what looks like a Bayesian posterior, as clear from Theorem 2 in the Appendix.

4.4. The introduction of censoring

To make the model more realistic, assume there are legal provisions on the market such that there exists a maximum recovery time $T^{\max}$, above which all recovery is exogenously stopped, notwithstanding the possibility to further recover later on. From a statistical point of view, the existence of $T^{\max}$ introduces the problem of right-censoring (Kleinbaum and Klein, 2012). In looking at real data, it means that we have no idea about what happens after $T^{\max}$: an exposure could have reached full recovery or not, but that remains unknown to us.

[3] We are using the simplified version of rule A, otherwise it is sufficient to set $q((1,m),c_i)=(0,0)$ for every i.

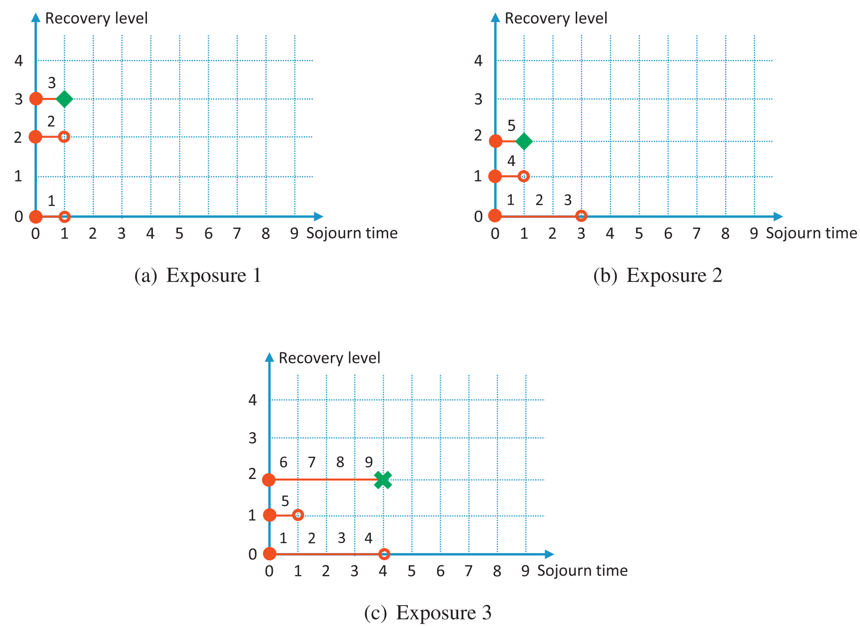

Fig. 1. The graphical representation of the recovery trajectories of exposures 1, 2, and 3 with maximum recovery time $T^{\max}=9$. For 1 and 2, the recovery process ends at times 3 and 5 (green diamonds, no censoring), and the passing of time can be easily read on the recovery trajectories. Notice that counterparty 2 never reaches the maximum recovery level $m=3$, but it stops in (1,2) with a write-off. For 3, the recovery process is right-censored at time $T^{\max}=9$ (green cross, censoring).

As an example, set $T^{\max}=9$, so that the recovery process of a given counterparty cannot last more than 9 time units. Fig. 1 gives a graphical representation of the recovery processes of exposures 1, 2 and 3 of Section 4.3 in the case of censoring. For counterparty 1 and 2, the recovery process ends at time $T_1=1+1+1=3<9$ and $T_2=3+1+1=5<9$, respectively, well before the maximum time limit, therefore no censoring takes place. For exposure 3, on the contrary, $T_3=4+1+6=11>9$, hence its recovery is forcedly interrupted and right-censored at time $T^{\max}=9$. The recovery trajectory for counterparty 3 in case of censoring is thus given by

$$R_3^{cens}=\{X_{18},\dots,X_{29}\}=\{(0,0),(1,0),(2,0),(3,0),(4,0),(0,1),(1,1),(0,2),(1,2),(2,2),(3,2),(4+,2)\},$$

where $(4+,2)$ indicates that the time spent in level 2 is censored in 4, given that the global $T^{\max}$ has been reached. Again, in case of censoring, we do not know what happens to the recovery trajectory of exposure 3 after $T^{\max}$.

For every counterparty i, we will make use of a dummy variable ρ_i to indicate whether its recovery trajectory is censored (1) or not (0). More in the Appendix.

4.5. Main properties

The R-RUP is characterized by several probabilistic properties, which make it powerful and flexible, and able to update its a priori knowledge with the information coming from actual recovery trajectories, even when these are censored, as very common in real business life (Resti and Sironi, 2007).

When recurrent, the R-RUP may be represented as a mixture of semi-Markov chains, and this mixture is characterizable as a new bivariate random distribution given by a particular interaction of Dirichlet and beta-Stacy processes (Theorem 1). The knowledge of this random distribution allows for the Bayesian prediction of the possible recovery trajectories of new counterparties (Theorem 2 and Corollary 1).

All the theorems, the proofs and the mathematical details related to the R-RUP are collected in the Appendix at the end of the paper.

5. Application: recovery modeling of family loans with US data

We show how the R-RUP can be used to jointly model recovery times and recovery levels, using US data from Freddie Mac (2017). The results we get are encouraging, as they show the ability of the R-RUP to learn from data and to reach interesting out-of-sample performances, also being able to correct possible mistakes in the elicitation of our prior beliefs.

5.1. Data and model initialization

We use a subset of the larger Single Family Loan-Level Dataset of Freddie Mac (2017), freely available online[4], covering around 23 million fixed-rate mortgages originated between January 1, 1999 and March 31, 2016. For each loan several covariates are available, the most relevant ones being the loan size, the loan-to-value, the FICO score (as indicator of the borrower's creditworthiness) and several measures of credit performance. For each single exposure the monthly EAD is registered together with all the repayments and, in case of a default, the information about the recovered amounts, the losses, the recovery procedure and the possible REO (real estate owned) dispositions is available.

Using Freddie Mac classification, a loan experiencing more than 180 days of delinquency is considered defaulted and the credit event is registered. From that very moment the recovery process starts, therefore delinquency day 180 corresponds to state (0,0) in the R-RUP construction. It is important to stress that we do not take into consideration those loans that have been liquidated prior to a delinquency of 180 days because of a short sale, a repurchase or a foreclosure. This choice is due to the fact that we need a clear definition of credit event to correctly define the initial state of our reinforced urn process. Consistently with the literature (Altman et al., 2005a), we also exclude all loans that, despite being formally defaulted after 180 days, have cured after the credit event and then prepaid or repurchased.

[4] http://www.freddiemac.com/news/finance/sf_loanlevel_dataset.html.

Table 1
Repartitions of the training set (in-sample: 20113 loans originated in the first quarter of 2006) and of the validation set (out-of-sample: 20038 loans originated in the second quarter of 2006) using the FICO scores, and the size of the loans (in thousand dollars).

FICO score	$\leq$ 650	(650,685]	(685,725]	$>$ 725	
Training	4691	4936	5057	5509	
Validation	4851	4747	4999	5441	
Loan size	$\leq$ 100	(100,150]	(150,200]	(200,250]	$>$ 250
Training	3346	4386	4333	3198	4850
Validation	3317	4608	4407	3111	4595

To train our model, that is to update the initial urn compositions (more in Section 5.2 and the Appendix), we start with a set of 20,113 defaulted loans originated in the first quarter of 2006. These loans correspond to roughly 10% of all the loans originated in the months of January, February and March 2006. Given our definition of credit event, these loans start defaulting in the third quarter of 2006, but many of them later. Almost all loans in the Freddie Mac Dataset are not right-censored: a complete recovery process is indeed observed, probably thanks to the long time window over which the loans are followed. As a consequence, in what follows, $\rho_i = 0$ for $i = 1, \ldots, 20113$, where ρ_i is the dummy variable indicating censoring.

In our training set almost 93% of the defaulted loans have experienced a write-off, and only 7% a full recovery.

The performances of the model have then been tested in-sample and, most importantly, out-of-sample, making inference and prediction about the recovery times and the recovery levels of loans originated in subsequent periods, like the 20,038 defaulted loans generated in the second quarter of 2006.

The 20,113 loans of the training sample and the 20,038 of the validation sample have been divided into different groups, to better model their recovery processes. In particular, we have considered 4 classes in terms of FICO score: ≤ 650, (650,685], (685,725], >725; and 5 classes in terms of size of the exposure: ≤ 100, (100,150], (150,200], (200,250] and >250 in thousand dollars. Table 1 gives the number of observations in each class in the training and in the validation sets according to the two classifications.

Splitting the data into groups guarantees that the assumption of exchangeability can be reasonably made for the counterparties in the different classes. Within each group, we thus assume that all the defaulted exposures are homogeneous and exchangeable in terms of risk (McNeil et al., 2015). While it is plausible to assume that counterparties with similar FICO scores (or loan sizes) are exchangeable, it is not at all safe to assume that all loans together are. Each class will be modeled with a different R-RUP.

For each defaulted loan in the Freddie Mac Dataset the monthly recovered amounts are collected. By dividing these numbers by the corresponding Exposure-at-Default, we easily obtain the monthly recovery rates. In order to be used in our model, these recovery rates are discretized into recovery levels. We have defined 13 levels:

$$L = \left\{ \begin{matrix} 0: 0\%, \ 1: (0\%, 10\%), \ 2: [10\%, 20\%), \ 3: [20\%, 30\%), \ 4: [30\%, 40\%), \\ 5: [40\%, 50\%), \ 6: [50\%, 60\%), \ 7: [60\%, 70\%), \ 8: [70\%, 80\%), \\ 9: [80\%, 90\%), \ 10: [90\%, 100\%), \ 11: 100+\%, \ 12: \text{termination} \end{matrix} \right\}. \quad (4)$$

In level 11, corresponding to full recovery, we have also included the few cases (in the order of tens) of recovery rates above 100%, while level 0 includes the few (again in the order of tens) situations in which the actual recovery rate is slightly negative, because the defaulted exposure has generated fees that the debtor has not paid. For more details about the technicalities of these events we refer to Freddie Mac (2017).

Given the range of variation of the recovery times in the data, we have chosen $t = 0, 1, \ldots, 100$, where each time step represents 1 month. The R-RUP is thus representable as a 13×101 matrix of urns, with levels of recovery on the rows, and times on the columns.

5.2. Prior elicitation and update

A fundamental aspect of the R-RUP initialization is the definition of the starting urn compositions, which represent our prior beliefs about recovery times and recovery levels. To embed our a priori in the R-RUP, we need to intervene on the number of balls in each urn. At the end of the Appendix, we give all the details about the process of prior elicitation using the properties of the bivariate random distribution that characterizes the R-RUP. Here it is sufficient to say that we have used three different prior sets to test the model. In all sets, the priors on the risk levels are discrete uniforms, with the empirically-based assumption that the higher the recovery level an exposure reaches, the higher the chance of further recovery (Altman et al., 2005b). Regarding the recovery times, conversely, we use the empirical cumulative distribution function in Prior Set 1, a discrete uniform with empirically determined support in Prior Set 2, and a uniform on [0,100] in Prior Set 3.

As in all Bayesian models, the prior elicitation step is very important in the R-RUP, in that it influences the speed at which the model converges to the true posterior distribution, and thus all estimates and predictions. But there are good news: even a totally wrong prior can be corrected with a sufficient number of observations, provided it is not degenerate, in accordance to Cromwell's rule (Jackman, 2009).

The urn compositions and our beliefs are then updated using the 20,113 recovery trajectories of the training sample. This is done by translating the recovery process of each observed counterparty into a sequence of samplings from the urns of the R-RUP, according to the rule of motion we have described in Section 4. For instance, when in the data we observe that counterparty i moves from state (t, l) to state $(t+1, l)$, we read this as an extraction of a c_l ball from the urn in (t, l). Therefore the composition of that urn is changed by adding r balls of color c_l. And so on for all the transitions we observe in the data, one counterparty at a time, under the hypothesis of exchangeability.

The parameter r thus plays the role of learning parameter. The bigger r, the quicker the model learns and adapts to the empirical data. The smaller r, the longer it will take to update our a priori. The calibration of r is therefore one of the ways in which we can show how confident we are about our beliefs, and how much we accept to modify them. Once again the importance of r diminishes with the number of available observations. With $r > 0$, thousands of observations will always be able to modify our a priori, shifting it towards the empirical reality that emerges from data. In the following we choose four different values for the reinforcement, $r \in \{0, 0.01, 1, 100\}$, where $r = 0$ means that we do not update our a priori, as if we do not trust data, while $r = 100$ indicates that empirical evidence is able to quickly modify our prior beliefs. For more details, once again we refer to the Appendix.

A natural question then arises: why do we need to elicit any a priori if we will always converge towards the empirical distribution of the data, with a sufficient number of observations? The answer is that: first, our prior beliefs become extremely important when the number of observations is not large, as in low default portfolios for instance (BCBS, 2011); second, a correct prior may compensate for the lack of information in the data and the problems of historical bias, as in the case of extreme and rare events (Shackle, 1955, Taleb, 2007); third, with the right prior elicitation we can also embed meaningful ideas about not-yet-observed trends and future developments. In the limit, in the totally ideal case in which our pri-

Table 2
P-values for two sample Kolmogorov-Smirnov Goodness-of-Fit Tests for the total recovery times in the training sample (20113 defaulted exposures originated in the first quarter of 2006), comparing the posterior R-RUP distributions and the empirical ones, for different reinforcements r. The training sample has been divided into 5 groups in terms of loan size, as per Table 1.

Loan size	Prior Set	$r=100$	$r=1$	$r=0.01$	$r=0$
$\le$ 100 K	1	0.13	0.08	0.15	0.00
	2	0.16	0.09	0.06	0.00
	3	0.09	0.01	0.23	0.00
(100 K,150 K]	1	0.01	0.00	0.00	0.00
	2	0.01	0.00	0.04	0.00
	3	0.02	0.07	0.02	0.00
(150 K,200 K]	1	0.10	0.05	0.11	0.00
	2	0.09	0.03	0.06	0.00
	3	0.02	0.06	0.00	0.00
(200 K,250 K]	1	0.90	0.58	0.26	0.00
	2	0.92	0.45	0.47	0.00
	3	0.39	0.22	0.22	0.00
$>$ 250 K	1	0.31	0.86	0.64	0.00
	2	0.46	0.74	0.07	0.00
	3	0.60	0.75	0.68	0.00

Table 3
P-values for two sample Kolmogorov-Smirnov Goodness-of-Fit Tests for the total recovery times in the training sample (20113 defaulted exposures originated in the first quarter of 2006), comparing the posterior R-RUP distributions and the empirical ones, for different reinforcements r. The training sample has been divided into 4 groups in terms of FICO scores, as per Table 1.

FICO score	Prior Set	$r=100$	$r=1$	$r=0.01$	$r=0$
$\le$ 650	1	0.01	0.02	0.03	0.00
	2	0.06	0.04	0.02	0.00
	3	0.04	0.06	0.00	0.00
(650,685]	1	0.02	0.04	0.02	0.00
	2	0.07	0.28	0.02	0.00
	3	0.04	0.13	0.07	0.00
(685,725]	1	0.21	0.16	0.03	0.00
	2	0.08	0.26	0.05	0.00
	3	0.08	0.06	0.21	0.00
$>$ 725	1	0.03	0.09	0.06	0.00
	2	0.06	0.03	0.08	0.00
	3	0.03	0.17	0.06	0.00

ors were true, we would need not a single observation in order to get perfect predictions.

5.3. Results

We now discuss the performances of the different R-RUP models, when initialized with one of the Priors Sets of Subsection 5.2, trained using actual data like the training set in Table 1, and then used to obtain the posterior distributions of the total recovery times, and of the recovery levels, for the different classes of interest.

Following Gelman et al. (2004), we start by testing the performances of our model in-sample, comparing the posterior distributions with the empirical distributions of the training data. This comparison takes the name of posterior consistency check (Meng, 1994).

Tables 2 and 3 show the p-values of several Kolmogorov–Smirnov goodness-of-fit tests (KS-test), between the posteriors and the empirical distributions of the total recovery times in the training sample, for the different groups of Table 1. We show the KS-tests as they tend to be more conservative, hence in principle against our model. Compatible results hold using other tests like the χ^2, omitted for the sake of space.

Let us consider Table 2, where the p-values for the posterior consistency check are provided for the loan sizes. Looking at loan size class $(200\,\text{K}, 250\,\text{K}]$, the R-RUP successfully passes the KS-test for most combinations of Prior Sets and reinforcement $r > 0$, at the 95% or 99% confidence level. As anticipated, the relative irrelevance of the value of r is due to the large number of observations available, together with the rather regular behavior of this class of exposures, which help the R-RUP in quickly learning from data. Conversely, as expected, if we do not let the R-RUP learn from the data, setting $r = 0$, and we just compare our naive prior beliefs with the empirical distributions of the recovery times, the null hypothesis of same distribution is definitely rejected.

Interestingly, in the $(100\,\text{K}, 150\,\text{K}]$ class of Table 2, we need the strong reinforcement $r = 100$ in order not to reject the null hypothesis at the 1% significance level, indicating that this class probably contains more peculiar behaviors with respect to our prior beliefs (and the other classes), and the R-RUP thus requires a stronger reinforcement (or possibly more data) in order to deal with them. A check of the data confirms our suspects: in this class there is a very large variability in the total recovery times, almost twice the other classes. For all the other loan classes, results are in line with what we have just said.

In Table 3 the same type of analysis is performed using the FICO scores. In this case the R-RUP still shows interesting performances, and the best results are apparently obtained setting $r = 1$, that is a mild reinforcement, which averages our prior beliefs and empirical evidence. This is due to the fact that a strong reinforcement like $r = 100$ quickly amplifies recurring patterns in the data, so that the few unusual ones appear farther away from the bulk of the distribution, and this is known to have effects on the KS statistic, which is based on the supremum distance.

All in all, the R-RUP satisfactorily passes the posterior consistency check, when we deal with the total recovery times.

In Table 4 we show the p-values of the KS-tests for the posterior check when looking at the recovery levels per loan size. For the FICO score the results are completely comparable and they are omitted. The performances of the R-RUP on the recovery levels are extremely good, for most reinforcements, indicating that our prior beliefs are probably not too far from reality (Altman et al., 2005b), and only minor modifications are necessary. This holds also for the $(100\,\text{K}, 150\,\text{K}]$ class, which is no longer as problematic as for the recovery times. In a nutshell: all combinations of Prior Sets and positive reinforcements give very good results. This makes sense, if we remember that the prior sets mainly differ for the time part, while the levels part is always the same (Subsection 5.2).

Table 4
P-values for two sample Kolmogorov-Smirnov Goodness-of-Fit Tests for the recovery levels in the training sample (20113 defaulted exposures originated in the first quarter of 2006), comparing the posterior R-RUP distributions and the empirical ones, for different reinforcements r. The training sample has been divided into 5 groups in terms of loan size, as per Table 1.

Loan size	Prior Set	$r=100$	$r=1$	$r=0.01$	$r=0$
$\le$ 100 K	1	0.99	1.00	0.97	0.00
	2	1.00	0.98	0.98	0.00
	3	1.00	0.99	1.00	0.00
(100 K,150 K]	1	0.80	0.31	0.88	0.00
	2	0.82	0.80	0.72	0.00
	3	0.82	0.51	0.52	0.00
(150 K,200 K]	1	0.38	0.83	0.46	0.00
	2	0.49	0.32	0.25	0.00
	3	0.55	0.40	0.46	0.00
(200 K,250 K]	1	1.00	1.00	1.00	0.00
	2	1.00	1.00	1.00	0.00
	3	1.00	0.99	1.00	0.00
$>$ 250 K	1	0.99	1.00	0.97	0.00
	2	1.00	1.00	1.00	0.00
	3	1.00	0.99	1.00	0.00

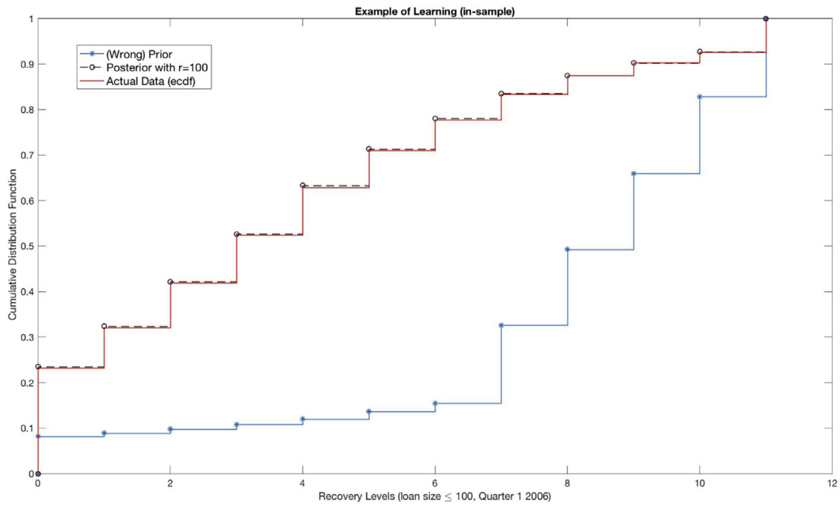

Fig. 2. Example of (in-sample) learning, showing the ecdf of the actual recovery levels in the first quarter 2016, a wrong prior and the posterior after the training of the model with $r = 100$. As expected, the model posterior very well approximates the actual distribution from the data.

Table 5
P-values for two sample Kolmogorov-Smirnov Goodness-of-Fit Tests between the posterior distributions of the trained R-RUP and the empirical distributions of the validation sample (20038 defaulted exposures originated in the second quarter of 2006) for the total recovery times in the 4 FICO score classes.

FICO score	Prior Set	$r = 100$	$r = 1$	$r = 0.01$	$r = 0$
$\leq$ 650	1	0.37	0.27	0.71	0.00
	2	0.79	0.24	0.99	0.00
	3	0.83	0.19	0.20	0.00
(650,685]	1	0.20	0.21	0.38	0.00
	2	0.48	0.10	0.21	0.00
	3	0.12	0.05	0.21	0.00
(685,725]	1	0.21	0.21	0.06	0.00
	2	0.03	0.17	0.26	0.00
	3	0.31	0.20	0.06	0.00
$>$ 725	1	0.54	0.69	0.46	0.00
	2	0.40	0.96	0.76	0.00
	3	0.60	0.86	0.30	0.00

Table 6
Median total recovery times per loan size class as predicted by the trained R-RUP against the actual values in the validation sample, for different configurations of priors and reinforcements r. Time expressed in months. Data in Table 1.

Loan size	Prior Set	Actual	$r = 100$	$r = 1$	$r = 0.01$	$r = 0$
$\leq$ 100 K	1	13	13	13	13	21
	2	13	13	13	13	18
	3	13	13	13	13	123
(100 K,150 K]	1	15	14	14	14	16
	2	15	14	14	15	9
	3	15	14	14	15	123
(150 K,200 K]	1	15	15	14	15	17
	2	15	14	15	15	9
	3	15	14	14	14	123
(200 K,250 K]	1	14	14	14	13	15
	2	14	14	13	13	9
	3	14	13	13	13	123
$>$ 250 K	1	13	13	13	13	17
	2	13	13	13	13	10
	3	13	13	13	13	124

As further evidence of the good in-sample performances of the R-RUP, in Fig. 2 we show the comparison between the empirical cumulative distribution function (ecdf) of the actual recovery levels in the 1st Quarter of 2016, for loan size class ≤ 100 (see Table 1), with a purposely wrongly elicited a priori, and with the updated posterior with reinforcement $r = 100$. As expected, given the good results with the KS tests, the trained model posterior very well approximates the actual distribution from the data, showing the ability of learning and improving, even when starting from a wrong set of beliefs.

Satisfied with the in-sample performances of the R-RUP, we can move to the more interesting out-of-sample validation. Is the trained R-RUP able to predict the recovery times and levels of future companies?

Following again Gelman et al. (2004), we can compare the posterior marginal distributions generated by the trained R-RUP (on the first quarter of 2006) with the empirical distributions of the validation sample, for example the one containing the 20,038 defaulted exposures originated in the second quarter of 2006. Theorem 2 in the Appendix proves extremely helpful in this situation. In Table 5, it is nice to see how the trained R-RUP is able to well approximate the distribution of the total recovery times in the validation sample. Similar results are obtainable using the loan size classes and/or the recovery levels, and they are available upon request.

Table 7
Median total recovery times per FICO scores class as predicted by the trained R-RUP against the actual values in the validation sample, for different configurations of priors and reinforcements r. Time expressed in months. Data in Table 1.

FICO score	Prior Set	Actual	$r = 100$	$r = 1$	$r = 0.01$	$r = 0$
$\leq$ 650	1	16	15	15	15	25
	2	16	15	15	15	15
	3	16	15	15	15	123
(650,685]	1	15	14	14	14	16
	2	15	14	14	14	9
	3	15	14	14	14	123
(685,725]	1	14	14	14	14	21
	2	14	14	14	13	15
	3	14	14	14	14	124
$>$ 725	1	13	12	13	12	17
	2	13	12	12	12	7
	3	13	13	13	12	124

Let us now consider some predictions made by the trained R-RUP with respect to the validation set, using the methodology we explain in the Appendix. Tables 6–9 compare the predicted median recovery times and levels for the defaulted exposures originated in the second quarter of 2006 with the actual medians from the validation set. The choice of the median as quantity of interest for credit risk purposes is consistent with works like that of

Table 8
Median recovery levels per loan size class as predicted by the trained R-RUP against the actual values in the validation sample, for different configurations of priors and reinforcements r. Levels follow the classification in Eq. (4). Data in Table 1.

Loan size	Prior Set	Actual	$r = 100$	$r = 1$	$r = 0.01$	$r = 0$
≤ 100 K	1	3	3	3	3	9
	2	3	3	3	3	10
	3	3	3	3	3	11
(100 K,150 K]	1	5	5	5	5	7
	2	5	5	5	5	10
	3	5	5	5	5	10
(150 K,200 K]	1	5	5	5	5	7
	2	5	5	5	5	10
	3	5	5	5	5	10
(200 K,250 K]	1	5	5	5	5	7
	2	5	5	5	5	11
	3	5	5	5	5	10
> 250 K	1	6	6	6	7	8
	2	6	6	6	7	10
	3	6	6	6	6	11

Table 9
Median recovery levels per FICO scores class as predicted by the trained R-RUP against the actual values in the validation sample, for different configurations of priors and reinforcements r. Levels follow the classification in Eq. (4). Data in Table 1.

FICO score	Prior Set	Actual	$r = 100$	$r = 1$	$r = 0.01$	$r = 0$
≤ 650	1	5	5	5	5	9
	2	5	5	5	5	11
	3	5	5	5	5	11
(650,685]	1	5	5	5	5	7
	2	5	5	5	5	11
	3	5	5	5	5	11
(685,725]	1	5	5	5	6	9
	2	5	5	5	6	10
	3	5	5	5	5	11
> 725	1	5	5	5	5	7
	2	5	5	5	7	11
	3	5	5	5	7	11

Table 10
P-values for two sample Kolmogorov–Smirnov Goodness-of-Fit Tests for the recovery levels, when comparing the posterior distributions of the R-RUP trained on 2006 data, with the empirical ones from the fourth quarter of 2009 (1252 data points), for different reinforcements r and the usual Prior Sets, according to the four FICO score classes.

FICO score	Prior Set	$r = 100$	$r = 1$	$r = 0.01$	$r = 0$
≤ 650	1	0.01	0.01	0.01	0.00
	2	0.01	0.03	0.00	0.00
	3	0.02	0.02	0.00	0.00
(650, 685]	1	0.00	0.00	0.00	0.00
	2	0.00	0.00	0.00	0.00
	3	0.00	0.00	0.00	0.00
(685, 725]	1	0.00	0.00	0.00	0.00
	2	0.00	0.00	0.00	0.00
	3	0.00	0.00	0.00	0.00
> 725	1	0.00	0.00	0.00	0.00
	2	0.00	0.00	0.00	0.00
	3	0.00	0.00	0.00	0.00

Table 11
P-values for two sample Kolmogorov-Smirnov Goodness-of-Fit Tests for the total recovery times, when comparing the posterior distributions of the R-RUP trained on loans originated in the first quarter of 2006 and the fourth quarters of 2007 and 2008, with the empirical ones from the fourth quarter of 2009, for different configurations of priors, reinforcements r, and FICO score classes.

FICO score	Prior Set	$r = 100$	$r = 1$	$r = 0.01$	$r = 0$
≤ 650	1	0.99	0.96	0.95	0.10
	2	0.96	0.97	0.72	0.00
	3	0.96	0.96	0.03	0.00
(650, 685]	1	0.22	0.14	0.17	0.00
	2	0.16	0.23	0.02	0.00
	3	0.14	0.17	0.00	0.00
(685,725]	1	0.27	0.19	0.22	0.00
	2	0.29	0.29	0.21	0.00
	3	0.24	0.21	0.00	0.00
> 725	1	0.04	0.01	0.00	0.00
	2	0.02	0.01	0.11	0.00
	3	0.00	0.00	0.00	0.00

Peluso et al. (2015), where medians are chosen for their statistical robustness.

Looking at Table 6, for example, we see that in the validation dataset, the actual median total recovery time is 15 months in loan size class (100 K, 150 K], and 13 months in class > 250 K. In both situations the R-RUP is able to provide predicted median values that are very close to the actual ones. The equivalence is also supported statistically by Mann-Whitney tests on medians. As before, the null reinforcement $r = 0$ is the one giving the worst results, and we find this comforting, as it underlines, once again, the ability of the R-RUP of learning and updating itself.

Upon request, we are glad to share the comparable results we obtain using other validation samples, like the defaulted exposures originated in the different quarters of 2007.

In line with expectations, bad performances are obtained if we move further into the future, still using only data from 2006 to train the model. The behaviors of the total recovery times and levels in 2008 and 2009 are not correctly reproduced, notwithstanding the r value and the Prior Set. For instance, in Table 10 we show the bad predictive power of the R-RUP trained on 2006 data, when dealing with the recovery levels of the defaulted exposures originated in the fourth quarter of 2009, according to the FICO classes. The only exception is represented by the class ≤ 650, consisting of the least reliable loans. This makes sense to us: being the worst class, the possibility of getting worse is limited, therefore it is more stable and easier to predict.

The decrease in the goodness of fit and in the predictive power for 2008 and 2009 is probably due to the impact of the 2008 economic crisis, which has substantially changed the dynamics of the recovery processes, given the larger number of defaults and the higher uncertainty on the markets (Caselli et al., 2008, Jankowitsch et al., 2014). To satisfactorily capture the dynamics of late 2008 or 2009, the R-RUP thus needs to be re-trained on 2007 and 2008 data.[5] Given its probabilist properties, and in particular conjugacy (Theorem 2), updating a R-RUP is rather simple, and it does not require to perform all computations anew. It is in fact sufficient to update its parameters using the new observations, as we show in the Appendix.

In Tables 11 and 12 we show the performances of the R-RUP for the total recovery times, if we complement its initial training on 2006 data with additional observations from 2007 and 2008. In particular, always focusing on the fourth quarter[6], we add the recovery trajectories of 20,118 defaulted exposures originated in 2007, and 3900 trajectories from 2008. We then use the re-trained model to make predictions with respect to the 1252 defaulted exposures originated in the fourth quarter of 2009. The p-values in the tables clearly show the ability of the R-RUP to correctly predict the total recovery times. Interestingly, we can observe relatively poorer performances in two cases: class > 725 for FICO scores,

[5] Or better priors have to be used. In a little experiment, on the basis of what we know happened in 2008 and 2009, we have modified our priors, and the R-RUP actually works well, because the a priori compensate for the lack of empirical information.

[6] Similar results hold for the other quarters, and combinations of quarters.

Table 12
P-values for two sample Kolmogorov–Smirnov Goodness-of-Fit Tests for the total recovery times, when comparing the posterior distributions of the R-RUP trained on loans originated in the first quarter of 2006 and the fourth quarters of 2007 and 2008, with the empirical ones from the fourth quarter of 2009, for different configurations of priors, reinforcements *r*, and loan size classes.

Loan size	Prior Set	$r = 100$	$r = 1$	$r = 0.01$	$r = 0$
$\leq$ 100 K	1	0.02	0.02	0.00	0.00
	2	0.03	0.01	0.01	0.00
	3	0.01	0.01	0.00	0.00
(100 K,150 K]	1	0.17	0.12	0.18	0.00
	2	0.18	0.18	0.13	0.00
	3	0.15	0.17	0.00	0.00
(150K,200K]	1	0.15	0.12	0.15	0.00
	2	0.16	0.19	0.14	0.00
	3	0.17	0.10	0.00	0.00
(200 K,250 K]	1	0.75	0.77	0.73	0.03
	2	0.79	0.82	0.02	0.00
	3	0.77	0.78	0.00	0.00
$>$ 250 K	1	0.41	0.46	0.30	0.00
	2	0.38	0.40	0.93	0.00
	3	0.39	0.42	0.00	0.00

Table 13
Median recovery levels as predicted by the trained R-RUP, using data from 2006 to 2008, against the actual values in the validation sample (4th quarter of 2009), for different configurations of priors, reinforcements *r*, and loan size classes.

Loan size	Prior Set	Actual	$r = 100$	$r = 1$	$r = 0.01$	$r = 0$
$\leq$ 100 K	1	6	4	4	4	7
	2	6	4	4	4	10
	3	6	4	4	4	11
(100 K,150 K]	1	7	5	5	5	8
	2	7	5	5	6	10
	3	7	5	5	6	11
(150 K,200 K]	1	7	6	6	6	7
	2	7	6	6	6	10
	3	7	6	6	6	10
(200 K,250 K]	1	8	6	6	6	7
	2	8	6	6	7	11
	3	8	6	6	7	11
$>$ 250 K	1	8	7	7	7	8
	2	8	7	7	7	10
	3	8	7	7	7	10

Table 14
P-values for two sample Kolmogorov–Smirnov Goodness-of-Fit Tests for the total recovery times, when comparing the posterior distributions of the R-RUP trained on yearly data from 2008, with the empirical ones from 2009, for different configurations of priors, reinforcements *r*, and FICO score classes.

FICO score	Prior Set	$r = 100$	$r = 1$	$r = 0.01$	$r = 0$
$\leq$ 650	1	0.93	0.94	0.92	0.66
	2	0.95	0.75	0.00	0.00
	3	0.95	0.97	0.00	0.00
(650,685]	1	0.77	0.76	0.89	0.97
	2	0.74	0.63	0.00	0.00
	3	0.79	0.89	0.00	0.00
(685, 725]	1	0.97	0.94	0.87	0.78
	2	0.96	0.99	0.00	0.00
	3	0.97	0.83	0.00	0.00
$>$ 725	1	0.88	0.89	0.98	0.19
	2	0.93	0.66	0.00	0.00
	3	0.80	0.97	0.00	0.00

Table 15
P-values for two sample Kolmogorov–Smirnov Goodness-of-Fit Tests for the total recovery times, when comparing the posterior distributions of the R-RUP trained on yearly data from 2008, with the empirical ones from 2009, for different configurations of priors, reinforcements *r*, and loan size classes.

Loan size	Prior Set	$r = 100$	$r = 1$	$r = 0.01$	$r = 0$
$\leq$ 100 K	1	0.96	0.95	0.96	0.66
	2	0.96	0.77	0.00	0.00
	3	0.97	0.99	0.00	0.00
(100 K, 150 K]	1	0.38	0.48	0.49	0.52
	2	0.48	0.27	0.00	0.00
	3	0.42	0.51	0.00	0.00
(150 K,200 K]	1	0.84	0.80	0.70	0.29
	2	0.88	0.95	0.00	0.00
	3	0.92	0.65	0.00	0.00
(200 K, 250 K]	1	0.57	0.59	0.42	0.31
	2	0.61	0.77	0.00	0.00
	3	0.59	0.33	0.00	0.00
$>$ 250 K	1	0.32	0.34	0.42	0.87
	2	0.36	0.37	0.00	0.00
	3	0.34	0.32	0.00	0.00

and class ≤ 100 in terms of loan size. Our explanation is that these two classes are those historically experiencing lower default rates, higher recovery levels and shorter recovery times. A sudden change in the recovery trajectories, as that observed during the crisis, is therefore more difficult to grasp. Nevertheless we consider the overall results quite satisfactory. Again, notice that as expected the naive Prior Set 3 is the one with the worst performances.

If we move from recovery times to recovery levels, we get similarly good results: the re-trained R-RUP is able to improve its predictive power, obtaining interesting performances, well above those of the old R-RUP trained on 2006 data only. All tables, omitted for the sake of space, are available upon request, also for the subsequent years up to the end of 2015. Here we only show Table 13, which contains the actual and the predicted median recovery levels in the 4th quarter of 2009, using data from 2006, 2007 and 2008. It is interesting to notice how the R-RUP, in this specific case, slightly underestimates the ultimate recovery level, on average by one level, showing a conservative behavior[7] The reason are once again unusual observations at the end of 2009, and, as before, our already good predictions could be improved by inputing more data in the R-RUP, for instance using those from the first half of 2009. It goes without saying that better priors, i.e. better experts' judgements, or a different scale (finer for instance) for the recovery levels could also be viable solutions.

The good results we obtain for quarters also hold if we consider yearly observations. In Tables 14 and 15 we show the comparisons at the yearly level for the recovery times, using the data up to the end of 2008 to predict the times in 2009. Once again the R-RUP performs well, learning from the data, even after a major economic crisis. Please notice that, in order to lower the computational burden, for the yearly comparisons we have used a smaller dataset, always provided by Freddie Mac.[8]

All in all, the R-RUP construction is clearly able to model recovery trajectories in a satisfactory way, being capable of adapting to periods of crisis, by constantly updating its performances. And that is actually the idea: to use the model under continuous reinforcement, every time new data become available. As said, updating the R-RUP does not require to perform all computations anew, but only to add the new information.

[7] Given the recent discussions about the margins of conservatism, and the preference of regulators for conservative estimates, slightly underestimating the actual recovery level would be something acceptable in the Basel framework (EBA, 2017). Unfortunately, we cannot guarantee that the R-RUP is always conservative, for all parameter choices. For sure a conservative prior set could be used as a standard, and even proposed by regulators.

[8] Citing Freddie Mac (2017): "The sample dataset is a simple random sample of 50,000 loans selected from each full vintage year and a proportionate number of loans from each partial vintage year of the full Single Family Loan-Level Dataset."

It is important to stress that our R-RUP has been trained using a lot of observations, which naturally correct wrong priors via reinforcement. In case of smaller datasets, the reliability of the prior beliefs becomes fundamental in order to obtain satisfactory results. In case of no prior knowledge, as common in Bayesian statistics (Dey and Rao, 2005), we naturally suggest the use of the empirical distribution of the training set as starting point. If used in the IRB setting (BCBS, 2011), specific prior sets could also be imposed by the regulator, which could also define restrictions for the reinforcement parameter.

5.4. The impact of discretization

By construction, the R-RUP relies on the discretization of recovery times and rates. While the discretization of time is not a problem, given that observations are often taken at pre-determined time points and, as discussed in Altman et al. (2005a), discrete time may even be an advantage in dealing with recovery risk, it may be worth analyzing the impact of discretizing recovery rates into recovery levels.

In Section 3, we said that one needs to define $m+1$ recovery levels, where level 0 accounts for no recovery, levels 1 to $m-1$ represent intermediate stages of recovery up to full recovery, and level m is the official termination level guaranteeing the technical condition of recurrence (Appendix: Lemma 1). The larger m, the finer the partition for the discretized recovery rates.

In the application to the Freddie Mac data, we have used the recovery levels of Eq. (4). In that partition, each level between 1 and 10 accounted for an extra 10% recovery. Do our results change if we modify the interpretation of each level, or if we chose a different number of levels?

Let us consider the following three alternatives:

$$L_1=\left\{\begin{array}{l}0:0\%,\ 1:(0\%,2\%),\ 2:[2\%,18\%),\ 3:[18\%,28\%),\ 4:[28\%,36\%),\\ 5:[36\%,44\%),\ 6:[44\%,51\%),\ 7:[51\%,60\%),\ 8:[60\%,71\%),\\ 9:[71\%,88\%),\ 10:[88\%,100\%),\ 11:100+\%,\ 12:\text{termination}\end{array}\right\},$$

$$L_2=\left\{\begin{array}{l}0:0\%,\ 1:(0\%,h\%),\ 2:[h\%,2h\%),\ 3:[2h\%,3h\%),\ 4:[3h\%,4h\%),\\ 5:[4h\%,5h\%),\ 6:[5h\%,6h\%),\ 7:[6h\%,7h\%),\ 8:[7h\%,8h\%),\\ 9:[8h\%,9h\%),\ 10:[9h\%,10h\%),\ 11:[10h\%,11h\%),\ 12:[11h\%,100\%),\\ 13:100+\%,\ 14:\text{termination}\end{array}\right\},$$

where $h=8.\bar{3}$, and

$$L_3=\left\{\begin{array}{l}0:0\%,\ 1:(0\%,12.5\%),\ 2:[12.5\%,25\%),\ 3:[25\%,37.5\%),\ 4:[37.5\%,50\%),\\ 5:[50\%,62.5\%),\ 6:[62.5\%,75\%),\ 7:[75\%,87.5\%),\ 8:[87.5\%,100\%),\\ 9:100+\%,\ 10:\text{termination}\end{array}\right\}.$$

In L_1 we use the same number of levels of L in Eq. (4), but we change their interpretation: they are no longer equally-spaced in terms of recovery. The intervals are empirically chosen so that each of them contains the same amount of observations: about one-tenth. In L_2 we increase the number of levels to 14, with a finer partition, in which each level between 1 and 12 accounts for an extra recovery of $8.\bar{3}$%. Finally, in L_3 we decrease the number levels, and each intermediate level accounts for an extra recovery of 12.5%.

Table 16 contains the same information of Table 2, when we substitute L with L_1. Since the intervals in L_1 are chosen empirically, to guarantee the same number of observations per level, this partition optimizes the information in the dataset in terms of number of updates per level, improving the goodness-of-fit (but not changing the conclusions substantially). Table 17 shows the results of the goodness-of-fit for L_2: in this case no dramatic difference is observed with respect to Table 2, apart from a general worsening for $r=0.01$. The same holds when using L_3.

The choice of L_1 slightly improves also the results for the goodness-of-fit of the ultimate recovery rate, but not for the predictive medians of times and levels. Conversely, L_2 and L_3 do not seem to have any particular influence.

Table 16
P-values for two sample Kolmogorov-Smirnov Goodness-of-Fit Tests for the total recovery times in the basic training sample (20113 defaulted exposures originated in the first quarter of 2006), comparing the posterior R-RUP distributions and the empirical ones, for different reinforcements r, under L_1.

Loan size	Prior Set	$r=100$	$r=1$	$r=0.01$	$r=0$
$\leq$ 100 K	1	0.84	0.76	0.86	0.00
	2	0.78	0.66	0.08	0.00
	3	0.89	0.88	0.20	0.00
(100 K, 150 K]	1	0.85	0.97	0.64	0.00
	2	0.48	0.96	0.02	0.00
	3	0.39	0.91	0.05	0.00
(150K, 200K]	1	0.92	0.30	0.62	0.00
	2	0.36	0.96	0.05	0.00
	3	0.97	1.00	0.05	0.00
(200 K, 250 K]	1	1.00	0.91	0.96	0.00
	2	0.96	0.54	0.06	0.00
	3	0.83	0.98	0.07	0.00
> 250 K	1	0.58	0.86	0.94	0.00
	2	0.87	0.99	0.08	0.00
	3	0.74	0.93	0.45	0.00

Table 17
P-values for two sample Kolmogorov–Smirnov Goodness-of-Fit Tests for the total recovery times in the basic training sample (20113 defaulted exposures originated in the first quarter of 2006), comparing the posterior R-RUP distributions and the empirical ones, for different reinforcements r, under L_2.

Loan size	Prior Set	$r=100$	$r=1$	$r=0.01$	$r=0$
$\leq$ 100 K	1	0.13	0.20	0.23	0.00
	2	0.51	0.09	0.00	0.00
	3	0.49	0.18	0.00	0.00
(100 K,150 K]	1	0.07	0.02	0.03	0.00
	2	0.01	0.00	0.00	0.00
	3	0.02	0.05	0.00	0.00
(150 K,200 K]	1	0.19	0.23	0.36	0.00
	2	0.36	0.07	0.00	0.00
	3	0.11	0.23	0.00	0.00
(200 K,250 K]	1	0.28	0.70	0.68	0.00
	2	0.55	0.45	0.00	0.00
	3	0.29	0.62	0.00	0.00
> 250 K	1	0.77	0.37	0.55	0.00
	2	0.37	0.67	0.00	0.00
	3	0.90	0.53	0.00	0.00

In our application, the R-RUP shows to be quite robust with respect to alternative choices of the recovery level partition. This is probably due to the large number of observations available, and–to some extent–to the fact that the partitions used above are not extremely different. For equally-spaced partitions with $m>25$, results are more problematic, given the increasing number of intervals with just a few observations.

In general, choosing a smaller number of levels improves fitting, because the number of observations and transitions per level increases, reinforcing the Bayesian learning process, ceteris paribus. Similarly, as seen above, an improvement is observed when intervals are chosen to guarantee more or less the same number of observations per level. Conversely, increasing m too much may lead to the situation in which for a specific level no transition is observed, so that our a priori is not changed, and, if our beliefs are wrong (or not meant to compensate an alleged lack of information in the data), this has naturally an impact on the goodness of fit, notwithstanding the choice of the reinforcement r.

As a natural trade-off, changing m has also effects in terms of precision. In the limit, we could consider a partition with just 4 levels, in which level 2 corresponds to a recovery rate in the interval (0%, 100%). This would dramatically increase the goodness of fit, but we all agree it would be useless, as we would end up saying that most counterparties reach some recovery rate between

0% and 100%, without distinction. On the opposite side we could be very precise and have thousands of levels, each representing an additional recovery of just a few decimals, but our a priori would be almost never updated for many levels (apart from the difficulty of eliciting it), and the definition of the strength of the reinforcement would become extremely complicated.

The best way to decide the number of levels is to find a compromise between precision, as required by internal procedures or the regulator, and the quantity and the quality of the empirical data. The more and the better the observations, the more precise the partition can be, but in any case each level should be characterized by a minimum number of transitions in order to exploit the Bayesian learning mechanism. To improve fitting, rarely visited levels should be aggregated, or, at least, the corresponding reinforcement calibrated to maximize the available information. Once again, experts' judgements could represent a viable solution.

6. Conclusions

In this paper we have introduced a new nonparametric survival approach for the modeling of recovery rates and recovery times. We have made use of the reinforced urn process (RUP) construction of Muliere et al. (2000) to build what we call the Recovery RUP (R-RUP).

The new process exhibits many interesting probabilistic properties, like partial exchangeability, semi-Markovianity and conjugacy. Its construction is intuitive and allows for simple simulations. The Bayesian update mechanism embedded in the Pólya urns builds a model that is able to combine prior knowledge, for example in the form of expert judgements, with empirical evidence; a model that learns and adapts to new phenomena and trends emerging from data, even in the case of censoring. The rate of updating can be controlled by acting on the strength of the a priori, and on the reinforcement mechanism of the single Pólya urns.

A first application on the Freddie Mac Single Family Loan-Level Dataset shows that the R-RUP actually provides interesting performances in terms of Bayesian prediction, with results that could be easily improved, if we were able to elicit more reliable priors, capable of compensating the possible lack of evidence in the data (the so-called historical bias) or unusual behaviors in certain periods.

In our approach we have not dealt with wrong-way risk, another interesting component of recovery risk, that is the positive dependence between PD and LGD observed in the empirical literature, especially during financial crises (Johnston Ross and Shibut, 2015, Dermine and De, 2006, Altman et al., 2005b). For us, in fact, default is given and it corresponds to the beginning of the recovery process, so that the modeling of PD is not relevant. An extension of the model to include wrong-way risk is in our future plans, and interesting starting points are the works of Han (2017) and Bade et al. (2011).

Acknowledgements

Pasquale Cirillo acknowledges the support of the EU Marie Sklodowska-Curie Career Integration Grant “Multivariate Shocks” (PCIG13-GA-2013-618794).

Both authors are thankful to the editor and two anonymous referees for their comments, which helped a lot in improving the quality of the work.

Appendix

We here collect all the mathematics related to the R-RUP construction, including definitions, theorems, proofs and technical details. Please notice that, not to overload notation, vector quantities are not expressed in bold.

Notation, the process $\{Y_p\}$ *and recurrence*

A process is said recurrent when, in infinite time, it will visit a given state infinitely often with probability one. For us this point is the origin (0,0).

From Muliere et al. (2000) we know that a recurrent RUP is partially exchangeable in the sense of Diaconis and Freedman (1980), defining a sequence of visited states that can be collected into exchangeable blocks, each one representing a Markov chain. Therefore, a recurrent RUP gives rise to a mixture of Markov chains with a given de Finetti (mixing) measure (Muliere et al., 2000).

Can we obtain something similar for the R-RUP? The answer is yes.

Fix $\eta_0^0 = 0$, and let the random variable $\eta_i^0 = \inf\{j > \eta_{i-1}^0 : X_j = (0,0)\}$ represent the i-th time the R-RUP process visits the point (0,0) for all nonnegative integer i, defining what Muliere et al. (2000) call a 0-block, i.e. a sequence of states in S starting with (0,0) and containing no further (0,0).[9] Thus $R_i = \{X_{\eta_{i-1}^0}, \dots, X_{\eta_i^0 - 1}\}$ will be the recovery history, with all the intermediate stages of recovery, of the ith defaulted counterparty, until its recovery process stops, because of a write-off, a full recovery or censoring. Let $\{R_i\}_{i=1}^k$ be the sequence of successive 0-blocks in $\{X_n\}$, representing the recovery histories of the first k defaulted exposures.

Let ψ be a mapping projecting all the finite sequences of elements of S, starting with an initial state (0,0) and ending up with (1, m), with no (0,0) appearing in between, into a sequence of recovery levels belonging to L. For $k \in \mathbb{N}_0$, any $t_0, \dots, t_k \in \mathbb{N}_0$ and $l_1, l_2, \dots, l_k \in L$, we have

$$\psi((0,0), \dots, (t_0, 0), \dots, (0, l_k), \dots, (t_k, l_k), (0, m), (1, m)) = \left\{ \overbrace{0, \dots, 0}^{t_0 \text{ times}}, \overbrace{l_1, \dots, l_1}^{t_1 \text{ times}}, \dots, \overbrace{l_k, \dots, l_k}^{t_k \text{ times}}, m \right\}.$$

Notice that $\psi(R_k)$ is just another way of representing the recovery process in block R_k, just focusing on the recovery levels.

The one-to-one correspondence between R_k and $\psi(R_k)$ is revealed by a simple example with termination level $m = 4$. For a block

$$R_k = \{(0,0), (1,0), (0,2), (1,2), (0,3), (1,3), (0,4), (1,4)\},$$

we have $\psi(R_k) = (0, 2, 3, 4)$. In this new representation, to recognize a 0-block, we have to look for the first 0 following termination level m.

Since the map ψ is measurable and bijective, and, in case of recurrence, the blocks $\{R_k\}_{k \geq 1}$ are exchangeable, the sequence $\{\psi(R_k)\}_{k \geq 1}$ is also exchangeable (Diaconis and Freedman, 1980), and we can use the sequence of levels $\{\psi(R_k)\}$ to build a new process $\{Y_p\}_{p \geq 0}$.

Assume that the first two recovery trajectories we observe are

$$R_1 = \{(0,0), (1,0), (0,2), (1,2), (0,3), (1,3), (0,4), (1,4)\}$$

and

$$R_2 = \{(0,0), (1,0), (2,0), (3,0), (0,1), (1,1), (0,2), (1,2), (0,4), (1,4)\}.$$

Then, the corresponding realization of the process $\{Y_p\}$ is

$$\{Y_0, Y_1, Y_2, Y_3, Y_4, Y_5, Y_6, Y_7, Y_8, Y_9\} = \{\psi(R_1), \psi(R_2)\}$$

[9] Notice that, given our assumptions on the urns, every 0-block necessarily ends with the state (1, m).

$$= \{0, 2, 3, 4, 0, 0, 0, 1, 2, 4\}.$$

We define all the notations we need to use afterwards. For any $l \in L$, set $\xi_0^l = 0$, and for the positive integer n, define recursively

$$\xi_n^l = \inf\{t > \xi_{n-1}^l : Y_{t-1} \neq l,\ Y_t = l\},$$

as the nth time the process $\{Y_p\}_{p\geq 0}$ touches l. For every $l \in L$, define a sequence $\{\tau_n^l\}_{n\geq 1}$ as follows:

$$\tau_n^l = \inf\{\, t - \xi_{n-1}^l : t > \xi_{n-1}^l,\ Y_t \neq l\}. \tag{5}$$

If $\{X_n\}$ is recurrent, for any $l \in L$ and $t \geq 0$, define

$$\delta_n^l = Y_{\xi_{n-1}^l + \tau_n^l}.$$

All in all, τ_n^l summarizes the sojourn time at recovery level l during the n-th visit, while δ_n^l is the next recovery level touched by the R-RUP after the nth visit at level l.

Set the initial level and the initial jump time as $L_0 = J_0 = 0$. For any positive integer n, define then

$$J_n = \inf\{t > J_{n-1} : Y_t \neq Y_{t-1}\}.$$

Set also $L_n = Y_{J_n}$ and $M_n = J_n - J_{n-1}$. Clearly J_n is the nth jump time of process $\{Y_p\}$, L_n is the level the process jumps to at time J_n, and M_n is the sojourn time at level L_{n-1} before jumping to level L_n.

From a probabilistic point of view, a necessary and sufficient condition for the R-RUP to be recurrent, so that state (0,0) is visited infinitely many times, is provided by the following Lemma.

Lemma 1 (Recurrence of the R-RUP)**.** *The R-RUP $\{X_n\}$ is recurrent if and only if, for all $l \in L$,*

$$\lim_{n\to\infty} \prod_{t=0}^{n} \frac{N_{(t,l)}(c_l)}{\sum_{i\geq l} N_{(t,l)}(c_i)} = 0, \tag{6}$$

where $N_{(t,\ l)}(c_i)$ is the number of balls of color $c_i \in C$ in urn $U((t,\ l))$.

Proof. From Lemma 2.13 and Lemma 3.23 in Muliere et al. (2000), it is clear that Eq. (6) is a necessary condition. Then, to prove that the process is recurrent, it is sufficient to show

$$\mathbb{P}\left[\bigcap_{u=1}^{\infty}\left\{\eta_u^0 < \infty\right\}\right] = 1.$$

First, $\mathbb{P}\left[\eta_1^0 < \infty\right] = 1$ holds since

$$\mathbb{P}[\eta_1^0 < \infty] \geq \mathbb{P}\left[\bigcap_{l=0}^{m}\{\tau_1^l < \infty\}\right] = 1, \tag{7}$$

where the equality is obtained from Lemma 3.23 in Muliere et al. (2000) and the finiteness of the maximum level m. Then, by induction, we can prove on n that

$$\mathbb{P}\left[\bigcap_{u=1}^{n+1}\left\{\eta_u^0 < \infty\right\}\right] = 1,$$

if $\mathbb{P}\left[\bigcap_{u=1}^{n}\left\{\eta_u^0 < \infty\right\}\right] = 1$. Since

$$\mathbb{P}[\eta_{n+1}^0 < \infty] \geq \mathbb{P}\left[\bigcap_{l=0}^{m}\bigcap_{u=1}^{n+1}\{\tau_u^l < \infty\}\right] = 1,$$

we get that

$$\mathbb{P}\left[\bigcap_{u=1}^{n+1}\left\{\eta_u^0 < \infty\right\}\right] = 1,$$

and the result follows. □

Lemma 1plays with the rule of motion q, by requiring that the probability for the R-RUP to stay at a given level l for infinite time is zero. This forces the R-RUP starting in (0,0) to jump to higher levels, reaching either full recovery (sampling of a c_m ball at recovery level $m-1$) or write-off (sampling of a c_m ball at recovery level $l < m-1$), before visiting the termination level m for one time unit and then restarting from (0,0). Clearly, when $\{X_n\}$ is recurrent, $\{Y_p\}$ is recurrent as well.

In what follows, all the definitions, propositions, lemmas, and theorems are based on the assumption of recurrence of the process $\{X_n\}$, or equivalently of the process $\{Y_p\}$.

The R-RUP as mixture of semi-Markov chains

Let us give some important definitions.

Definition 1 (beta-Stacy Process)**.** The random distribution function F is a beta-Stacy process with jumps at $t \in \mathbb{N}_0$ and parameters $\{\alpha_t, \beta_t\}_{t\in\mathbb{N}_0}$, if there exist mutually independent random variables $\{V_t\}_{t\in\mathbb{N}_0}$, each beta distributed with parameters $(\alpha_t,\ \beta_t)$, such that the random mass assigned by F to $\{t\}$, written $F(\{t\})$, is given by $V_t \prod_{u<t}(1 - V_u)$.

Introduced by Walker and Muliere (1997), the beta-Stacy process can be seen as a generalization of the well-known Dirichlet process, a pivotal random distribution in Bayesian nonparametrics (Dey and Rao, 2005). Its RUP representation was first given in Muliere et al. (2000).

Definition 2 (beta-Stacy Dirichlet Process)**.** The random probability mass function Q on $L \times \mathbb{N}_0$ is called beta-Stacy Dirichlet (BSD) process with parameters $\{\alpha_t,\ \beta_t\}_{t\geq 0}$ and $\{\gamma_t\}_{t\geq 0}$, if there exist mutually independent Dirichlet processes $\{W_t\}_{t\geq 0}$ of parameters $\{\gamma_t\}$, and a beta-Stacy process F with jumps at $t \in \mathbb{N}_0$ and parameters $\{\alpha_t, \beta_t\}_{t\in\mathbb{N}_0}$, independent of the sequence $\{W_t\}$, such that the random mass assigned to the element $(j,\ t)$ in $L \times \mathbb{N}_0$ is $Q(j, t) = F(\{t\})W_t(j)$.

The BSD process is the bivariate random distribution that characterizes the R-RUP, as stated by Theorem 1 below. However, before considering that fundamental result, we first need to introduce the concept of semi-Markov chain.

Definition 3. A process $\{Y_p\}_{p\geq 0}$ is a discrete time semi-Markov chain on $L \times L \times \mathbb{N}_0$, if the values Y_{J_n} at its jump times form a Markov chain. Moreover, conditionally on Y_{J_n} and all the previous information, the sojourn time in state Y_{J_n}, and the next state the process jumps to, only depend on Y_{J_n}.

If $\{Y_p\}_{p\geq 0}$ is a discrete time semi-Markov chain, let $G_l(j,\ t)$ be the probability of staying at level l for a time t and then jumping to another state j, i.e.

$$G_l(j, t) = \mathbb{P}[L_{n+1} = j, M_{n+1} = t \mid L_n = l].$$

Then G is the semi-Markov kernel of $\{Y_p\}$ on $L \times L \times \mathbb{N}_0$.

Consider the set of semi-Markov kernels $\mathscr{G}$ on $L \times L \times \mathbb{N}_0$ in the topology of coordinate convergence. If the process $\{Y_p\}$ is a mixture of semi-Markov chains, there exists a probability measure κ, also called the mixing measure, on the Borel subset of the space $\mathscr{G}$, such that, for any $l_1, \dots, l_n \in L$, $t_1, \dots, t_n \geq 0$, $n \geq 1$ and $l_0 = 0$, we have

$$\mathbb{P}[(L_1, M_1) = (l_1, t_1), \dots, (L_n, M_n) = (l_n, t_n)] \\ = \int_{\mathscr{G}} \prod_{u=1}^{n} G_{l_{u-1}}(l_u, t_u)\kappa(\mathrm{d}G). \tag{8}$$

Proposition 1. *For every $l \in L$, the sub-sequence $\{(\delta_n^l, \tau_n^l)\}$ is exchangeable. Moreover, the sub-sequences*

$$\{(\delta_n^0, \tau_n^0)\}, \{(\delta_n^1, \tau_n^1)\}, \dots, \{(\delta_n^{m-1}, \tau_n^{m-1})\}, \{(\delta_n^m, \tau_n^m)\} \tag{9}$$

are mutually independent.

Proof. The exchangeability of the sequence $\{(\delta_n^l, \tau_n^l)\}$ simply derives from the fact that both δ_n^l and τ_n^l are measurable functions of the 0-blocks, which are exchangeable by construction. The independence, conversely, simply derives from the R-RUP construction.

□

The following theorem shows that the R-RUP $\{Y_p\}$ is a mixture of semi-Markov chains, i.e. the blocks $\psi(R_i)$ constituting $\{Y_p\}_{p\geq 0}$ are exchangeable and each of them is a semi-Markov chain. The mixing measure κ can be obtained by characterizing G. Knowing the functional form of the semi-Markov kernel G will prove essential to use the R-RUP in practice.

Theorem 1. *If recurrent, the process $\{Y_p\}_{p\geq 0}$ is a mixture of semi-Markov chains, and its mixing measure κ is unique. More precisely, there exists a unique random kernel G such that, conditionally on G, Y_p is a semi-Markov chain with semi-Markov kernel G.*

For each level of risk $l \in L$, the kernel G can be explicitly characterized as the product $G_l(j,t) = W_{(t,l)}(j) F^l(\{t\})$, for any $l, j \in L$ and $t \in \mathbb{N}_0$, where:

- *F^l is a beta-Stacy process with jumps at $t \in \mathbb{N}_0$, and parameters*

$$\alpha_t^l = \sum_{h\neq l} \frac{N_{(t,l)}(c_h)}{r} \quad \text{and} \quad \beta_t^l = \frac{N_{(t,l)}(c_l)}{r},$$

  with r the reinforcement of the different Pólya urns.
- *$\{W_{(t,l)}\}_{t\in\mathbb{N}_0, l\in L}$ are mutually independent Dirichlet processes of parameter $\gamma_{(t,l)}(\cdot)$, all independent of $\mathbf{F} = \{F^l, l\in L\}$, assigning mass $\frac{N_{(t,l)}(c_j)}{r}$ to the j-th component for $l+1 \leq j \leq m$, and mass 0 for $0 \leq j \leq l$.*

Proof. Muliere et al. (2000) show that a recurrent RUP $\{X_n\}$ defines a mixture of Markov chains. This means that there exists a random transition matrix R, which is a random element of the set of all stochastic matrices, such that for all finite sequences $(s_0, s_1, \ldots, s_n)$ of elements of S,

$$\mathbb{P}[X_0 = s_0, X_1 = s_1, \ldots, X_n = s_n \mid R] = \prod_{u=1}^{n} R(s_{u-1}, s_u) a.s. \tag{10}$$

Theorem 2.16 in Muliere et al. (2000) characterizes R by showing that its rows are mutually independent random probability masses on S, and for all $s \in S$, $R(s)$ follows a Dirichlet process with parameter $\omega(s)$, which assigns $\frac{N_s(c)}{r}$ to state $q(s, c) \in S$ with $c \in C$.

For any $l, j \in L$ and $t \in \mathbb{N}_0$, define

$$\begin{aligned} W_{(t,l)}(j) &= \frac{R[(t,l),(0,j)]}{1 - R[(t,l),(t+1,l)]}, \\ V_t^l &= 1 - R[(t,l),(t+1,l)], \\ F^l(\{t\}) &= V_t^l \textstyle\prod_{u<t}(1 - V_u^l). \end{aligned}$$

Because of the tail free property of the Dirichlet process (Dey and Rao, 2005),

$$W_{(t,l)}(\cdot) = \left[W_{(t,l)}(0), \ldots, W_{(t,l)}(m)\right]$$

follows a Dirichlet process with measure $\gamma_{(t,l)}(\cdot)$ assigning $N_{(t,l)}(c_j)/r$ to $j > l$ and 0 to the rest. Moreover, $\{V_t^l\}_{t\geq 0}$ are a series of independent beta distributed random variables with parameters (α_t^l, β_t^l) where $\alpha_t^l = \sum_{h\neq l} N_{(t,l)}(c_h)/r$ and $\beta_t^l = N_{(t,l)}(c_l)/r$.

From Walker and Muliere (1997), under the recurrence condition of Lemma 1, F^l is a discrete-time beta-Stacy process with parameters $\{\alpha_t^l, \beta_t^l\}$. This said, the goal is to characterize the mixing measure of the component $\{(\delta_n^l, \tau_n^l)\}$ of the semi-Markov chain $\{Y_p\}$.

According to the de Finetti representation theorem, the exchangeability of the sub-sequence $\{(\delta_n^l, \tau_n^l)\}$ guarantees the existence and uniqueness of a random probability measure Q_l on $L \times \mathbb{N}_0$, conditionally on which the elements of the sub-sequence are i.i.d. with a bivariate mass function Q_l (de Finetti, 1975). The following lemma characterizes this de Finetti measure as a new combination of well-known random distributions, namely the Dirichlet and the beta-Stacy processes. □

Lemma 2. *If recurrent, for any $l \in L$, the unique random probability measure Q_l on $L \times \mathbb{N}_0$ is a BSD process with parameters $\{\alpha_t^l, \beta_t^l\}_{t\geq 0}$ and $\{\gamma_{(t,l)}\}_{t\geq 0}$, where, for any $t \in \mathbb{N}_0$, $\alpha_t^l = \sum_{h\neq l} \frac{N_{(t,l)}(c_h)}{r}$, $\beta_t^l = \frac{N_{(t,l)}(c_l)}{r}$, and $\gamma_{(t,l)}(\cdot)$ assigns mass $\frac{N_{(t,l)}(c_j)}{r}$ to the jth component for $l < j \leq m$, and mass 0 to all the other components.*

Proof. For all $t_1, \ldots, t_n \in \mathbb{N}_0$ and $l_1, \ldots, l_n \in L$, we have

$$\begin{aligned} &\mathbb{P}[\delta_1^l = l_1, \tau_1^l = t_1, \delta_2^l = l_2, \tau_2^l = t_2, \ldots, \delta_n^l = l_n, \tau_n^l = t_n \mid R] \\ =\ & \textstyle\prod_{k=1}^n \prod_{u<t_k} R[(u,l),(u+1,l)] \times R[(t_k,l),(0,l_k)]) \\ =\ & \textstyle\prod_{k=1}^n \prod_{u<t_k} R[(u,l),(u+1,l)] \times \\ & (1 - R[(t_k,l),(t_k+1,l)]) \times \frac{R[(t_k,l),(0,l_k)]}{1 - R[(t_k,l),(t_k+1,l)]} \\ =\ & \textstyle\prod_{k=1}^n [W_{(t_k,l)}(l_k) V_{t_k}^l \prod_{u<t_k}(1 - V_u^l)] \\ =\ & \textstyle\prod_{k=1}^n [W_{(t_k,l)}(l_k) F^l(\{t_k\})] \quad a.s. \end{aligned}$$

Define $Q_l(j,t) = W_{(t,l)}(j) F^l(\{t\})$, for any $l, j \in L$ and $t \in \mathbb{N}_0$. Since the bivariate measure is measurable with respect to the $\mathbb{P}$-completion of the Borel σ-algebra of R, the above relation remains valid after replacing R with Q_l. Hence, for any $l_1, \ldots, l_n \in L$ and $t_1, \ldots, t_n \in \mathbb{N}_0$,

$$\mathbb{P}[\delta_1^l = l_1, \tau_1^l = t_1, \delta_2^l = l_2, \tau_2^l = t_2, \ldots, \delta_n^l = l_n, \tau_n^l = t_n \mid Q_l] = \prod_{k=1}^n Q_l(l_k, t_k) \ a.s.$$

The result follows. □

Lemma 3. *$F^0, \ldots, F^m$ are mutually independent. The elements $\{W_{(t,l)}\}_{t\in\mathbb{N}_0, l\in L}$ are mutually independent, and independent from $\mathbf{F} = \{F^l, l \in L\}$.*

Proof. It is sufficient to show that for any $l \in L$ and any $t, u \in \mathbb{N}_0$, $W_{(t,l)}$ and $F^l(\{u\})$ are independent. We have already seen that for every $l \in L$ and $u \in \mathbb{N}_0$, $\{W_{(t,l)} : t \neq u, t \in \mathbb{N}_0\}$ and V_u^l are independent thanks to the R-RUP construction. We only need to prove that $W_{(t,l)}$ and V_t^l are independent, which can be obtained from the tail free property of the Dirichlet process. The result then follows.

Define $G_l(\cdot,\cdot) = Q_l$, for any $l \in L$. Thanks to Lemmas 2 and 3, for any $l_1, \ldots, l_n \in L$ and $t_1, \ldots, t_n \in \mathbb{N}_0$, we have that

$$\mathbb{P}[(L_1, M_1) = (l_1, t_1), \ldots, (L_n, M_n) = (l_n, t_n) \mid G] = \prod_{u=1}^{n} G_{l_{u-1}}(l_u, t_u) \quad a.s. \tag{11}$$

Theorem 1 is thus finally proved. □

The posterior predictive distribution with and without right-censoring

The results we have obtained with recurrence and semi-Markovianity are extremely useful to derive the posterior predictive distribution of the R-RUP in terms of the BSD process, which proves to be conjugate (Dey and Rao, 2005).

Here below we give an important theorem, in which we show that, given the information about the recovery processes of k counterparties, the recovery trajectory of the $k+1$th one is still a semi-Markov chain with updated kernel. With updated kernel we indicate the one resulting from the combination of the prior knowledge, embedded into the initial composition of the urns of the R-

RUP, and the empirical evidence we can collect from data, i.e. the k observed recovery trajectories.

The result we provide deals with the general case of possibly censored observations, as per Subsection 4.4. Naturally it is completely valid even when all recovery trajectories are fully observed, being this just a special case.

Let $\rho_i,\ i=1,\dots,k,$ be a dummy variable indicating censoring. Each recovery block $R_i,\ i=1,\dots,k,$ can then be accompanied by a corresponding ρ_i, so that $(R_i, \rho_i=0)$ indicates that the recovery trajectory of counterparty i is not censored, while $(R_i, \rho_i=1)$ tells us that the recovery has been censored because of some $T^{\max}$.

The information in $(R_i,\ \rho_i)$ can be further summarized by $H_i=(\boldsymbol{\psi}(R_i),\rho_i)$, where H_i represents the recovery history of counterparty $i=1,\dots,k$, including the information about censoring. With $\boldsymbol{H_k}=[H_1,H_2,\dots,H_k]$ we indicate the histories of the first k counterparties in our portfolio. Let $s_{(t,\,l)}(j)$ be the number of counterparties sojourning at level l in t and then jumping to another level $j\in L$, w_t^l be the number of counterparties whose exact sojourning time at level l exceeds time t, and v_t^l be the number of counterparties whose sojourning time at level l is right-censored in t.

Theorem 2. *Conditionally on $\boldsymbol{H_k}$, possibly with right-censoring, the element H_{k+1}, representing the recovery history of the $(k+1)$-th counterparty, is (still) a semi-Markov chain with semi-Markov kernel $\hat{G}$, such that $\hat{G}_l(j,t)=E\big[\tilde{G}_l(j,t)\big]$, for $l,\ j\in L$ and $t\in\mathbb{N}_0$, where $\tilde{G}_l(j,t)=\tilde{F}^l(\{t\})\tilde{W}_{(t,l)}(j)$ and*

- *$\tilde{F}^l$ is a beta-Stacy process with jumps at $t\in\mathbb{N}_0$ and updated parameters*
  $$\tilde{\alpha}_t^l = \alpha_t^l + \textstyle\sum_{h>l} s_{(t,l)}(h),\qquad \tilde{\beta}_t^l = \beta_t^l + w_t^l + v_t^l.$$
- *$\{\tilde{W}_{(t,l)}\}$ are mutually independent Dirichlet with parameters $\tilde{\gamma}_{(t,l)}$ where*
  $$\tilde{\gamma}_{(t,l)}(j) = \gamma_{(t,l)}(j) + s_{(t,l)}(j),$$
  for $t\in\mathbb{N}_0$ and $j\in L$.
  As in Theorem 1, $\{\tilde{W}_{(t,l)}\}$ is independent from $\tilde{F}^l$ for every $l\in L$.

Proof. Let $\rho_i^l, i=1,\dots,k_l$ be a binary variable indicating local right-censoring, where k_l is the number of counterparties sojourning at recovery level l. The recovery histories of the first k defaulted counterparties $\boldsymbol{H_k}$ can be split into several sets $\boldsymbol{H_k^l},\ l\in L$, each of which contains the observed sojourning time τ_i^l at level l, and the next recovery level δ_i^l, when $\rho_i^l=0$, or just the observed sojourning time if censoring has occurred with $\rho_i^l=1$. □

Lemma 4. *For every $l\in L$, given $\boldsymbol{H_k^l}$, the posterior predictive probability mass function for the couple $(\delta_{k_l+1}^l,\tau_{k_l+1}^l)$ is given by*

$$\mathbb{P}\left[\delta_{k_l+1}^l=j,\tau_{k_l+1}^l=t\mid \boldsymbol{H_k^l}\right]=\frac{\tilde{\alpha}_t^l}{\tilde{\alpha}_t^l+\tilde{\beta}_t^l}\prod_{u<t}\frac{\tilde{\beta}_u^l}{\tilde{\alpha}_u^l+\tilde{\beta}_u^l}\times\frac{\tilde{\gamma}_{(t,l)}(j)}{\tilde{\alpha}_t^l},$$

for $j\in L$ and all nonnegative integer t, with

$$\begin{aligned}\tilde{\alpha}_t^l &= \alpha_t^l + v_t^l,\\ \tilde{\beta}_t^l &= \beta_t^l + w_t^l,\\ \tilde{\gamma}_{(t,l)}(j) &= \gamma_{(t,l)}(j) + s_{(t,l)}(j),\end{aligned}\qquad (12)$$

where $s_{(t,\,l)}(j)$ is the number of counterparties sojourning at level l in t and then jumping to level j, where w_t^l is the number of counterparties whose exact sojourning time at level l exceeds time t, and where v_t^l is the number of counterparties whose sojourning time at level l is right-censored in t.

Proof. If right-censoring occurs, from Lemma 2, we have that,

$$\mathbb{P}[\tau_1^l=t_1,\rho_1^l=1\mid G]=1-\textstyle\sum_{u\le t_1}F^l(\{u\})=1-F^l(t_1).$$

Otherwise,

$$\mathbb{P}[\delta_1^l=l_1,\tau_1^l=t_1,\rho_1^l=0\mid G]=W_{(t_1,l)}(l_1)F^l(\{t_1\}).$$

Hence,

$$\begin{aligned}
&\mathbb{P}\left[(\tau_{k_l+1}^l,\delta_{k_l+1}^l)=(t,j)\mid \boldsymbol{H_k^l}\right]\\
&=\frac{\mathbb{P}[(\tau_{k_l+1}^l,\delta_{k_l+1}^l)=(t,j),\boldsymbol{H_k^l}]}{\mathbb{P}[\boldsymbol{H_k^l}]}\\
&=\frac{E\left[W_{(t,l)}(j)F^l(\{t\})\prod_{i=1}^{k_l}\left[W_{(t_i,l)}(l_i)F^l(\{t_i\})\mathbb{1}[\rho_i^l=0]+[1-F^l(t_i)]\mathbb{1}[\rho_i^l=1]\right]\right]}{E\left[\prod_{i=1}^{k_l}\left[W_{(t_i,l)}(l_i)F^l(\{t_i\})\mathbb{1}[\rho_i^l=0]+[1-F^l(t_i)]\mathbb{1}[\rho_i^l=1]\right]\right]}\\
&=\frac{E\left[W_{(t,l)}(j)F^l(\{t\})\prod_{\rho_i^l=0}\left[W_{(t_i,l)}(l_i)F^l(\{t_i\})\right]\prod_{\rho_i^l=1}\left[1-F^l(t_i)\right]\right]}{E\left[\prod_{\rho_i^l=0}\left[W_{(t_i,l)}(l_i)F^l(\{t_i\})\right]\prod_{\rho_i^l=1}\left[1-F^l(t_i)\right]\right]}\\
&=\frac{E\left[W_{(t,l)}(j)\prod_{\rho_i^l=0}W_{(t_i,l)}(l_i)\right]E\left[F^l(\{t\})\prod_{\rho_i^l=0}F^l(\{t_i\})\prod_{\rho_i^l=1}\left[1-F^l(t_i)\right]\right]}{E\left[\prod_{\rho_i^l=0}W_{(t_i,l)}(l_i)\right]E\left[\prod_{\rho_i^l=0}F^l(\{t_i\})\prod_{\rho_i^l=1}\left[1-F^l(t_i)\right]\right]}\\
&=\frac{\tilde{\alpha}_t^l}{\tilde{\alpha}_t^l+\tilde{\beta}_t^l}\prod_{u<t}\frac{\tilde{\beta}_u^l}{\tilde{\alpha}_u^l+\tilde{\beta}_u^l}\times\frac{\tilde{\gamma}_{(t,l)}(j)}{\tilde{\alpha}_t^l},
\end{aligned}$$

with the second last equality coming from Lemma 3, and where, for $j\in L$ and all nonnegative integer t,

$$\begin{aligned}\tilde{\alpha}_t^l &= \alpha_t^l + v_t^l,\\ \tilde{\beta}_t^l &= \beta_t^l + w_t^l,\\ \tilde{\gamma}_{(t,l)}(j) &= \gamma_{(t,l)}(j) + s_{(t,l)}(j).\end{aligned}\qquad (13)$$

Thanks to Proposition 1, Theorem 2 is then proved. □

Corollary 1. *The $(k+n)$th recovery history H_{k+n} for any positive integer n, conditionally on $\boldsymbol{H_k}$, possibly with right-censoring, is a semi-Markov chain with a common semi-Markov kernel $\hat{G}$ defined as in Theorem 2.*

Exploiting Theorem 2 and Corollary 1, we can perform Bayesian prediction about the recovery trajectories (levels and times) of future counterparties, given our a priori–as expressed by the initial compositions of the urns in the R-RUP–and the collected empirical evidence in $\boldsymbol{H_k}$, with or without censoring. In fact, Theorem 2 tells us that the semi-Markov kernel governing the future recovery histories is the updated BSD process combining our a priori with the data, and whose parameters we know. In statistical terms, the BSD process is conjugate. The probability of every feasible future recovery history can therefore be computed explicitly, given the available information.

The more our priori is close to the truth, the more reliable our predictions from the very beginning, and the better we can overcome the problem of censoring, by compensating the lack of information in the data with our beliefs. In case of a wrong a priori, however, in a way similar to what happens with machine learning, a sufficient amount of recovery histories can compensate unrealistic beliefs, as we see in Section 5.

Prior Elicitation

Notice that G is a random kernel on $L\times L\times\mathbb{N}_0$ and for every $l\in L$, G_l is a bivariate random probability measures taking values in the space of bivariate discrete probability measures, which can be obtained as the product of beta-Stacy and Dirichlet processes. Therefore, if we want to center G on a given semi-Markov kernel $\bar{G}$, we can exploit the following relationship between the prior guess $\bar{G}$ and the initial composition of the Pólya urns constituting the R-RUP. Generalizing results in Muliere et al. (2000) and

Mezzetti et al. (2007), we have

$$E[G_l(j,t)] = \bar{G}_l(j,t) = \frac{\alpha_t^l}{\alpha_t^l + \beta_t^l} \prod_{u<t} \frac{\beta_u^l}{\alpha_u^l + \beta_u^l} \times \frac{\gamma_{(t,l)}(j)}{\alpha_t^l},$$

for $j,\ l \in L,\ j > l,\ t \in \mathbb{N}_0$, and where $\alpha_t^l = \sum_{h=l+1}^m \frac{N_{(t,l)}(c_h)}{r}$, $\beta_t^l = \frac{N_{(t,l)}(c_l)}{r}$, and $\gamma_{(t,l)}(j) = \frac{N_{(t,l)}(c_j)}{r}$.

From this we derive that, for every state $(t,\ l) \in S$, and every color c_j, with $j = 0, \dots, m$, the initial urn composition is given by

$$N_{(t,l)}(c_j) = \begin{cases} d_t^l \times \bar{G}_l(j,t), & \text{if } j > l; \\ d_t^l \times \left[1 - \sum_{k=0}^t \sum_{h>l} \bar{G}_l(h,k)\right], & \text{if } j = l; \\ 0 & \text{if } j < l. \end{cases} \quad (14)$$

From Definition 3, the semi Markov kernel $\bar{G}$ is controlled by $\bar{F}^l$, charactering the distribution of the sojourn times in level l, and $\bar{W}_{(t,l)}$, expressing the probability of jumping to other levels from level l, given the time t spent in l. Hence, to elicit prior guesses easily and flexibly, we can just specify the one-dimensional distributions $\bar{F}^l$ and $\bar{W}_{(t,l)}$, replacing $\bar{G}_l$ in Eq. (14) with their product.

The parameter $d_t^l > 0$ in Eq. (14) is what Muliere et al. (2000) call the strength of belief. It is a value representing how confident we are in our a priori. The larger $d_t^l > 0$, the more we are convinced that our prior beliefs are correct, so that the R-RUP will need more time and more observations, or a stronger reinforcement r, in order to correct them if wrong. In this paper we set $d_t^l = 1$ for all $t \in \mathbb{N}_0$ and $l \in L$.

Regarding the reinforcement parameter r, which in the R-RUP represents the strength of updating, i.e. how empirical observations affect and modify our a priori, we have chosen four different values to compare, i.e. $r \in \{0,\ 0.01,\ 1,\ 100\}$. A value of $r = 0.01$ says that we do not trust the empirical evidence too much, and that we prefer to stick to our a priori, only allowing for very slow updates. A value of $r = 100$, conversely, indicates that we trust the empirical evidence and we are ready to update our beliefs quickly. Naturally for $r = 0$ the process never learns from data, and our a priori is never changed.

The following list collects the prior beliefs we have elicited to test our model, for the different classes (size of the exposure and FICO score) described in Subsection 5.1:

Prior Set 1: For every $l \in L$ and $t \in \mathbb{N}_0$, $\bar{W}_{(t,l)}$ assigns mass 0 to $\{0, \dots, l\}$ and equal mass to $\{l+1, \dots, m\}$, while $\bar{F}^l$ is the empirical cumulative distribution function (ecdf) of the sojourn times, as obtainable from the data, for each recovery level l.

Prior Set 2: For every $l \in L$ and $t \in \mathbb{N}_0$, $\bar{W}_{(t,l)}$ assigns mass 0 to $\{0, \dots, l\}$ and equal mass to $\{l+1, \dots, m\}$, while $\bar{F}^l$ is a discrete uniform distribution on $\{0, 1, \dots, T_{emp}^l\}$, with T_{emp}^l indicating the maximum empirical sojourn time at level l as observable in the training data.

Prior Set 3: For every $l \in L$ and $t \in \mathbb{N}_0$, $\bar{W}_{(t,l)}$ assigns mass 0 to $\{0, \dots, l\}$ and equal mass to $\{l+1, \dots, m\}$, while $\bar{F}^l$ a discrete uniform distribution on $\{0, 1, \dots, 100\}$.

Remember that $\bar{W}_{(t,l)}$ essentially governs the probability of jumping to higher recovery levels from $(t,\ l)$, while $\bar{F}^l$ accounts for the probability of the permanence time at level l. Please notice that the a priori elicited by $\bar{W}_{(t,l)}$, with the support $\{l+1, \dots, m\}$ depending on the level l, corresponds to assuming that the higher the recovery level an exposure reaches, the higher the chance of further recovery. This is consistent with the empirical evidence about recovery rates in most countries (Altman et al., 2005a, Resti and Sironi, 2007).

Given Priors Sets 1–3, Eq. (14) is used to define the initial composition of all the urns of the R-RUPs in the 13×101 matrix of visitable states. Naturally, the urn compositions' limitations of Subsection 4.2 still apply.

References

Altman, E., Resti, A., Sironi, A., 2005a. Recovery Risk: The Next Challenge in Credit Risk Management. Risk Books.

Altman, E.I., Brady, B., Resti, A., Sironi, A., 2005b. The link between default and recovery rates: theory, empirical evidence, and implications. J. Bus. 78 (6), 2203–2227.

Bade, B., Rösch, D., Scheule, H., 2011. Default and recovery risk dependencies in a simple credit risk model. Eur. Financ. Manag. 17 (1), 120–144.

Bastos, J.A., 2010. Forecasting bank loans loss-given-default. J. Bank. Financ. 34 (10), 2510–2517.

BCBS, 2005. Guidance on Paragraph 468 of the Framework Document. July. URL: http://www.bis.org/publ/bcbs115.pdf.

BCBS, 2006. International Convergence of Capital Measurement and Capital Standards. June. URL: http://www.bis.org/publ/bcbs128.pdf.

BCBS, 2011. A global regulatory framework for more resilient banks and banking systems. June. URL: http://www.bis.org/publ/bcbs189.pdf.

Bellotti, T., Crook, J., 2012. Loss given default models incorporating macroeconomic variables for credit cards. Int. J. Forecast 28 (1), 171–182.

Bonini, S., Caivano, G., 2013. The survival analysis approach in basel ii credit risk management: modeling danger rates in the loss given default parameter. J. Credit Risk 9 (1), 101.

Calabrese, R., Zenga, M., 2010. Bank loan recovery rates: measuring and nonparametric density estimation. J. Bank. Financ. 34 (5), 903–911.

Caselli, S., Gatti, S., Querci, F., 2008. The sensitivity of the loss given default rate to systematic risk: new empirical evidence on bank loans. J. Financ. Serv. Res. 34 (1), 1–34.

Chawla, G., Forest Jr, L.R., Aguais, S.D., 2015. Point-in-time (pit) lgd and ead models for ifrs9/cecl and stress testing. J. Risk Manag. Financ. Inst. 9 (3), 249–263.

Cirillo, P., Hüsler, J., Muliere, P., 2010. A nonparametric urn-based approach to interacting failing systems with an application to credit resk modeling. Int. J. Theor. Appl. Financ. 41 (0), 1–18.

Derbyshire, J., 2017. The siren call of probability: dangers associated with using probability for consideration of the future. Futures 88, 43–54.

Dermine, J., De, C.N., 2006. Bank loan losses-given-default: a case study. J. Bank. Financ. 30 (4), 1219–1243.

Dey, D., Rao, C., 2005. Bayesian thinking, modeling and computation. Handbook of Statistics. Elsevier Science.

Diaconis, P., Freedman, D., 1980. De Finetti's theorem for markov chains. Ann. Probab. 8 (1), 115–130.

EBA, 2017. Guidelines on PD estimation, LGD estimation and the treatment of defaulted exposures. URL: https://www.eba.europa.eu/regulation-and-policy/credit-risk.

de Finetti, B., 1975. Theory of probability II. Wiley.

Freddie Mac, 2017. Freddie Mac single family loan-level dataset and user guide. URL: http://www.freddiemac.com/news/finance/sf_loanlevel_dataset.html.

Frye, J., 2005. The effects of systematic credit risk: a false sense of security. In: Recovery Risk. Riskbooks, London, pp. 187–200.

Gelman, A., Carlin, J.B., Stern, H.S., Rubin, D.B., 2004. Bayesian Data Analysis, 2nd ed. Chapman and Hall/CRC.

Gürtler, M., Hibbeln, M., 2013. Improvements in loss given default forecasts for bank loans. J. Bank. Financ. 37 (7), 2354–2366.

Han, C., 2017. Modeling severity risk under pd-lgd correlation. Eur. J. Financ. 23 (15), 1572–1588.

Hartmann-Wendels, T., Miller, P., Töws, E., 2014. Loss given default for leasing: parametric and nonparametric estimations. J. Bank. Financ. 40 (1), 364–375.

Höochstötter, M., Nazemi, A., 2013. Analysis of loss given default. Invest. Manag. Financ. Innovations 10 (4), 70–79.

Huang, X., Oosterlee, C.W., 2011. Generalized beta regression models for random loss given default. J. Credit Risk 7 (4), 45–70.

Hull, J.C., 2015. Risk management and financial institutions, 4th edition Wiley.

IFRS9, 2014. Financial Instruments. IFRS Foundation. URL: http://tinyurl.com/khrqs42

Im, J.-k., Apley, D.W., Qi, C., Shan, X., 2012. A time-dependent proportional hazards survival model for credit risk analysis. J. Oper. Res. Soc. 63 (3), 306–321.

Jackman, S., 2009. Bayesian Analysis for the Social Sciences. Wiley.

Jankowitsch, R., Nagler, F., Subrahmanyam, M.G., 2014. The determinants of recovery rates in the us corporate bond market. J Financ. Econ. 114 (1), 155–177.

Johnston Ross, E.B., Shibut, L., 2015. What drives loss given default? evidence from commercial real estate loans at failed banks. FDIC Center for Financial Research Paper. 2015-03

Jorion, P., 2009. Risk management lessons from the credit crisis. Eur. Financ. Manag. 15 (5), 923–933.

Khieu, H.D., Mullineaux, D.J., Yi, H.C., 2012. The determinants of bank loan recovery rates. J. Bank. Financ. 36 (4), 923–933.

Kiefer, N.M., 1988. Economic duration data and hazard functions. J. Econ. Lit. 26 (2), 646–679.

Kleinbaum, D.G., Klein, M., 2012. Survival Analysis. Springer.

Knight, F.H., 1964. Risk, Uncertainty and Profit (Reprint of the 1921 Edition). Augustus M. Kelley, Bookseller.

Krüger, S., Rösch, D., 2017. Downturn LGD modeling using quantile regression. J. Bank. Financ. 79, 42–56.
Loterman, G., Brown, I., Martens, D., Mues, C., Baesens, B., 2012. Benchmarking regression algorithms for loss given default modeling. Int. J. Forecast 28 (1), 161–170.
Mahmoud, H., 2008. Polya urn Models. CRC Press.
McNeil, A.J., Frey, R., Embrechts, P., 2015. Quantitative Risk Management: Concepts, Techniques and Tools. Princeton University Press, Princeton, NJ, USA.
Meng, X.-L., 1994. Multiple-imputation inferences with uncongenial sources of input. Statist. Sci. 9 (4), 538–558.
Mezzetti, M., Muliere, P., Bulla, P., 2007. An application of reinforced urn processes to determining maximum tolerated dose. Stat. Probab. Lett. 77 (7), 740–747.
Muliere, P., Secchi, P., Walker, S., 2000. Urn schemes and reinforced random walks. Stoch. Process. Their. Appl. 88 (1), 59–78.
Murphy, K.P., 2012. Machine Learning: A Probabilistic Perspective. The MIT Press.
Peluso, S., Mira, A., Muliere, P., 2015. Reinforced urn processes for credit risk models. J. Econ. 184 (1), 1–12.
Reitgruber, W., 2015. Methodological thoughts on expected loss estimation for ifrs 9 impairment: hidden reserves, cyclical loss predictions and lgd backtesting. Credit Technol. (92) 7–29.
Renault, O., Scaillet, O., 2004. On the way to recovery: a nonparametric bias free estimation of recovery rate densities. J. Bank. Financ. 28 (12), 2915–2931.
Resti, A., Sironi, A., 2007. Risk Management and Shareholders' Value in Banking. John Wiley & Sons.
Schmit, M., 2004. Credit risk in the leasing business-a case study of low probability of default. J. Bank. Financ. 28, 811–833.
Schuermann, T., 2004. What Do We Know About Loss Given Default?Credit Risk Models and Management, 1–32.
Shackle, G.L.S., 1955. Uncertainty in Economics and Other Reflections. Cambridge University Press.
Taleb, N.N., 2007. The Black Swan: The Impact of the Highly Improbable. Random House Group.
Walker, S., Muliere, P., 1997. Beta-stacy processes and a generalization of the Pólya-urn scheme. Ann. Stat. 25, 1762–1780.
Witzany, J., Rychnovský, M., Charamza, P., 2012. Survival analysis in LGD modeling. Eur. Financ. Account. J. 7 (1), 6–27.
Yashkir, O., Yashkir, Y., 2013. Loss given default modelling: comparative analysis. J. Risk Model Valid 7 (1), 25–59.
Zhang, J., Thomas, L.C., 2012. Comparisons of linear regression and survival analysis using single and mixture distributions approaches in modelling LGD. Int. J. Forecast 28 (1), 204–215.