

## RESEARCH ARTICLE

# Forecasting intermittent and sparse time series: A unified probabilistic framework via deep renewal processes

Ali Caner Türkmen<sup>1‡\*</sup>, Tim Januschowski<sup>1</sup>, Yuyang Wang<sup>2</sup>, Ali Taylan Cemgil<sup>3</sup>

**1** Amazon Web Services AI Labs, Berlin, Germany, **2** Amazon Web Services AI Labs, East Palo Alto, CA, United States of America, **3** Department of Computer Engineering, Boğaziçi University, Istanbul, Turkey

‡ Work done before joining Amazon.

\* [atturkm@amazon.com](mailto:atturkm@amazon.com)



## OPEN ACCESS

**Citation:** Türkmen AC, Januschowski T, Wang Y, Cemgil AT (2021) Forecasting intermittent and sparse time series: A unified probabilistic framework via deep renewal processes. PLoS ONE 16(11): e0259764. <https://doi.org/10.1371/journal.pone.0259764>

**Editor:** Yossiri Adulyasak, HEC Montréal, CANADA

**Received:** September 16, 2020

**Accepted:** October 26, 2021

**Published:** November 29, 2021

**Copyright:** © 2021 Türkmen et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

**Data Availability Statement:** The UCI data set is available on the UCI ML repository (<https://archive.ics.uci.edu/ml/index.php>). The Car Parts data set is available in the Monash Time Series Forecasting Repository (<https://zenodo.org/record/4656022#.YYPvWMY0-Ak>). The M5 data set is available on Kaggle (<https://www.kaggle.com/c/m5-forecasting-accuracy>). RAF and Auto data sets are available online at <https://github.com/canerturkmen/gluonts/tree/intermittent-datasets/datasets/>.

**Funding:** The first three authors of this article (ACT, TJ, and YW) are employed by Amazon Web

## Abstract

Intermittency are a common and challenging problem in demand forecasting. We introduce a new, unified framework for building probabilistic forecasting models for intermittent demand time series, which incorporates and allows to generalize existing methods in several directions. Our framework is based on extensions of well-established model-based methods to discrete-time renewal processes, which can parsimoniously account for patterns such as aging, clustering and quasi-periodicity in demand arrivals. The connection to discrete-time renewal processes allows not only for a principled extension of Croston-type models, but additionally for a natural inclusion of neural network based models—by replacing exponential smoothing with a recurrent neural network. We also demonstrate that modeling continuous-time demand arrivals, i.e., with a temporal point process, is possible via a trivial extension of our framework. This leads to more flexible modeling in scenarios where data of individual purchase orders are directly available with granular timestamps. Complementing this theoretical advancement, we demonstrate the efficacy of our framework for forecasting practice via an extensive empirical study on standard intermittent demand data sets, in which we report predictive accuracy in a variety of scenarios.

## 1 Introduction

Intermittent demand forecasting (IDF) is concerned with demand data where demand appears sporadically in time [1–4], i.e., long runs of zero demand are observed before periods with nonzero demand. Not only does this sparsity render most standard forecasting techniques impractical; it leads to challenges on measuring forecast accuracy [5], model selection [6], and forecast aggregation [7].

Demand for large shares of inventory catalogues in manufacturing are well known to exhibit intermittency [8, 9]. Intermittent demand is most likely to appear with slow-moving, (sometimes) high-value items that are critical to production processes. For example, spare parts in aerospace and defense are well known to exhibit intermittent patterns [10, 11].

Services. The funder provided support in the form of salaries for authors ACT, TJ, YW, but did not have any additional role in the study design, data collection and analysis, decision to publish, or preparation of the manuscript. The specific roles of these authors are articulated in the 'author contributions' section.

**Competing interests:** The first three authors of this article (ACT, TJ, and YW) are employed by Amazon Web Services. This does not alter our adherence to PLOS ONE policies on sharing data and materials.

Therefore, precise estimates of forecast uncertainty, *e.g.*, with access to forecast distributions, are vital for IDF.

IDF was recognized as a unique and challenging problem in the early 70s [1, 12–15]. Upon recognizing that traditional approaches such as simple exponential smoothing (SES) led to poor forecasts in slow-moving inventory, Croston [1] proposed to independently apply exponential smoothing to consecutive positive demand sizes, and to the number of periods between each (*i.e.*, *interarrival* or *interdemand* times). Towards developing forecast distributions and uncertainty estimates, statistical models underlying Croston's method were proposed by [16–19].

In this paper, we make a sequence of observations on IDF methods proposed so far. We start working from Croston's original insight, towards a consistent set of flexible intermittent demand models. We explore extensions of these models in several directions, which are natural consequences of our new framework to systematically tackle IDF. Our paper builds on early work in [20], expanding on theoretical concepts and adding a thorough empirical analysis. We summarize our observations and contributions below.

- We draw previously unrecognized connections between existing IDF models and renewal processes. We note that these two subjects of applied statistics both deal with temporal sparsity and have been developed for planning spare parts inventories.
- We introduce a flexible set of discrete-time renewal process models for stationary intermittent demand. We illustrate that these models are able to capture patterns such as temporal clustering, aging, and quasi-periodicity of demand.
- We cast Croston-like models as instances of discrete-time conditional (self-modulating) renewal processes. Through this observation, we can introduce recurrent neural networks to recover more flexible ways in which renewal processes are modulated.
- We observe that our construction applies in contexts in which individual demand events are observed in continuous time. We make connections to temporal point processes and their neural variants and apply these models in IDF.
- We evaluate our approach on both synthetic and real-world data sets. The experiments on synthetic data illustrate qualitatively and quantitatively how our proposed models can capture complications such as alternating periods better than prior methods. For the experiments on real-world data set, we report results on Car Parts, Auto, RAF, and M5 competition data sets, as well as a newly introduced dataset from the UCI repository. While our results are mixed, we can conclude that using probabilistic neural networks is a promising direction for further exploration, in particular for the data set from the M5 competition.

Our paper consists of four sections. In Section 2, we start with a survey of literature on intermittent demand forecasting, with a focus on machine learning-driven methodologies. In the subsequent Section 3, we introduce preliminary statistical and algorithmic concepts such as renewal processes, temporal point processes, and recurrent neural networks. In Section 4, we introduce our unified framework—which we test on a wide array of synthetic and real data experiments in the Section 5.

## 2 Literature review

The literature on IDF can be traced back to the original paper by Croston [1], and it has since grown considerably. In the following, we provide a chronological overview of the literature on IDF, with a focus on works that highlight a model-based approach or employ neural networks.

Croston [1] observed that more accurate point forecasts of intermittent demand resulted from performing exponential smoothing separately on sequences of demand sizes and inter demand intervals. Apart from providing a forecast heuristic, he introduced a set of models which, he posited, would lead to the forecast equations he provided. Later developed and corrected [21, 22], Croston's method has been the most common IDF method for stockists and forecasters [23], and the *de facto* standard in forecasting software and libraries. Among early extensions to Croston's method, Schultz [13] suggested applying different smoothing constants to interarrival times and demand. Willemain et al. [15] verified results by [1], who also illustrated that interdemand times exhibit autocorrelation by statistical analysis on IDF data.

Dunsmuir and Snyder [14] provided an explicit model for inventory control in an intermittent demand scenario. Notably, they took stochastic delivery lead times of replenishment to be gamma random variables, modeling replenishment in continuous time. Johnston and Boylan [16] pointed out that Croston's method did not offer an estimate of the variability in demand. They proposed separate continuous-time models, *i.e.*, temporal point processes, for demand sizes and order arrivals. Notably, they assumed the order arrivals to follow a Poisson process, building on Cox's results from renewal theory and developing estimates of demand mean and variance. They also explored different order size distributions under Poissonian order arrivals.

Syntetos [24] gave a survey of IDF in his thesis, along with error approximations in IDF models and some extensions for inventory control. Importantly, he pointed out the "inversion bias" in Croston's original forecast estimates and provided a revised estimator [22]. Syntetos later also gave a simpler approximant apart from reviewing questions around forecast accuracy measures in IDF [25]. A similarly modified IDF estimate was studied by Shale et al. [26], where the authors assumed underlying Poissonian order arrivals. A review of these studies was given in [10]. For the question of how to define intermittency, *i.e.*, which time series constitute intermittent demand we refer to [2, 3, 27–29].

Another method for generating point forecasts was given by Teunter et al. [30]. The authors applied exponential smoothing directly on the *probability* of demand, bypassing the inversion bias in Croston's method. A review, as well as a comparative study of point forecast methods on real IDF data was presented by Babai et al. [31], with no strong evidence in favor of any of the forecast functions considered.

The problem of associating a stochastic model to Croston's method was explored by Snyder [17]. He proposed several modifications to the method and considered intermittent demand forecasts via the parametric bootstrap on a set of state-space models (SSM). Notably, this work draws connections between the renewal-type model of Croston [1, Appx. B] and single source of error SSMs.

Shenstone and Hyndman [18] investigated the validity of Croston's methods on several fronts, pointing out that the i.i.d. assumption on order sizes and interarrival times were inconsistent with the exponentially weighted moving average (EWMA) forecast estimates. The authors explored a set of "modified" models that would yield Croston's estimates, assuming both sizes and intervals follow an ARIMA(0,1,1) process. To ensure nonnegativity, required of both sizes and intervals, they proposed log-Croston models where an ARIMA(0,1,1) process determined the required quantities in the log domain. Importantly, the authors argue that there exist no practicable models that yield Croston's estimates as optimal forecasts. This is since any local-level model that yields the EWMA as an unbiased one-step-ahead forecast, and is defined on the positive real half-line, suffers from a convergence problem [32, 33], *i.e.*, draws from the model converge to 0 over long terms. Let us note that the question of a stochastic model for IDF was also raised in [34, 35]. An extensive discussion on model-based IDF is given in [36, Ch. 16].

Snyder et al. [19] proposed an extended set of models, contrasting several ways in which stochastic Croston-like models were considered in the literature. They compared Poisson, negative binomial and zero-inflated Poisson distributions for demand sizes. For all distributions, they tested “static” (constant) parameters as well as undamped (IMA) and damped (AR) dynamic parameterizations. The proposed class of “dynamic” models are simple and are shown to outperform Croston’s method and static parameter models. In this paper, we connect previous model-based approaches in [18, 19] to the rich theory behind renewal processes. As such, for the first time, we can derive Croston’s methods and other proposed models from a principled probabilistic framework.

SSM approaches were considered and compared by Yelland [37], who also later proposed a hierarchical Bayesian treatment of the problem [38]. Seeger et al. [8] considered Bayesian SSM with multiple sources of error, focusing on applicability on very large inventory catalogues (see also [39] for more details). Their method incorporates exogenous features, and presents an efficient approximate inference algorithm enabling large-scale and distributed learning. Altay et al. [40] explored an adaptation of Croston’s method with Holt-like trends. Seasonality in IDF, with a comparison of Holt-Winters and SARIMA, was considered in [41]. Recently, a single source of error SSM with multiplicative errors was explored by Svetunkov and Boylan [42]. A novel method for forecast aggregation via temporal hierarchies, THieF, was introduced by Kourentzes and Athanasopoulos [43].

Using machine learning methods—mainly, neural networks—in forecasting is both a recently growing area of research and the subject of ongoing debate [44, 45]. In their extensive comparison of machine learning methods to traditional “statistical” methods, Makridakis et al. [46] find scarce evidence in favor of using neural networks or other ML-based methods in forecasting. The authors argue that most empirical evidence in forecasting favors “model-driven” methods over “data-driven,” in the terms of Januschowski et al. [47]. Moreover, data-driven methods are often harder to train and replicate; and require substantially greater effort for implementation and computation. The findings from this study are controversial to say the least as practical evidence from industrial researcher has consistently pointed to contrary conclusions (see e.g., [48–51]). Supporting this point of view are the results from highly visible M4 competition [52] which was won by a hybrid method that combined exponential smoothing, recurrent neural networks, and ensembling [53]. Other such hybrid approaches have appeared recently, mainly in the machine learning literature [48, 54–57], along with software libraries for neural network time series forecasting [58]. A thorough review of the use of neural networks in forecasting was recently given by Benidis et al. [59].

Several studies have considered neural networks in the context of IDF. Gutierrez et al. [60] experimented with a single hidden layer feedforward network, with only three hidden units, trained to predict the demand on the next time interval, given time since last demand and the last observed demand size. They reported favorable results compared to Croston and Syntetos-Boylan methods. These results were debated by Kourentzes [61], who compared several neural network architectures to a suite of naive and Croston-based forecasters and demonstrated low forecast accuracy in terms of mean absolute error (MAE). However, the author also found favorable performance by neural networks when inventory control metrics were considered directly. Mukhopadhyay et al. [62] experimented with several training strategies to find that neural networks outperform traditional methods. Recurrent neural networks were considered in [63], and several extensions were explored in [64]. A deep LSTM, in the context of IDF, appeared recently in [65]. Our paper allows combining neural networks with a model-based approach, a first in the literature for intermittent demand forecasting. As a consequence, the neural forecasting methods we present here are the first dedicated neural network-based models for IDF which yield probabilistic forecasts.

Finally, we remark that IDF requires special attention when evaluating forecasting accuracy as standard forecast accuracy metrics are well known to fail in this context. Hyndman [5] proposed MASE, mean absolute error scaled against a naive-one forecast, for the IDF task. Kim and Kim [66] advanced this notion to a reportedly more robust metric, employing trigonometric functions for comparing forecasts and ground truth. Although we used both of these metrics in our evaluations, we found both approaches to be potentially misleading, as a “zero forecast” often yielded better results in terms of MASE or MAAPE than any other method. As our methods are primarily geared towards probabilistic forecasting, we use P50 and P90 loss metrics as in [48]. Further, we also report root mean squared error (RMSE) and root mean squared scaled error (RMSSE), the main error metric used in the recent M5 competition [67].

### 3 Preliminaries

#### 3.1 Problem setup

We consider univariate, nonnegative, integer-valued time series, corresponding to random variables denoted  $\{Y_n\}_{n \in \{1, 2, \dots\}}$ ,  $Y_n \in \mathbb{N}$ . Here,  $n$  indexes uniformly spaced intervals in time, each corresponding to a *demand review period*. Realizations  $\{y_n\}$  of  $\{Y_n\}$  will typically contain long runs of zeros, *i.e.*, only rarely will  $y_n$  be greater than zero.

For a clearer exposition of our models, we will mostly use the *size-interval* notation for intermittent demand. As many  $Y_n$  are zero, it will be useful to define an auxiliary index  $i \in \mathbb{N}$ , such that  $i$  indexes the *issue points*, *i.e.*, periods where nonzero demand is observed. More formally, we define the one to one map  $\sigma(i) = \min\{n \mid \sum_{m=1}^n \mathbb{I}\{Y_m > 0\} \geq i\}$ , where  $\mathbb{I}\{.\}$  is the indicator function (Iverson bracket). We denote *interdemand times*—number of periods between issue points— $Q_i$ . That is,  $Q_i = \sigma(i) - \sigma(i - 1)$ , taking  $\sigma(0) = 0$ . We denote demand sizes for issue points  $M_i = Y_{\sigma(i)}$ . Random variables  $M_i$ ,  $Q_i$ , both defined on positive integers, fully determine  $Y_n$  and vice versa.

We use corresponding lowercase Latin letters to denote instantiations of random variables denoted in uppercase Latin.  $\mathbb{E}[\cdot]$  and  $\mathbb{V}[\cdot]$  denote mathematical expectation and variance respectively. Unless otherwise specified, we keep indices  $i$ ,  $n$  analogous to their function in this introduction, indexing issue periods and review periods respectively. Range indexes will denote consecutive subsequences of a random process or realization, *e.g.*,  $Y_{1:k} = \{Y_1, Y_2, \dots, Y_k\}$ . We reserve  $\mathcal{N}$ ,  $\mathcal{G}$ ,  $\mathcal{P}$ ,  $\mathcal{NB}$ ,  $\mathcal{E}$  to denote Gaussian, geometric, Poisson, negative binomial, and exponential distributions respectively. Unless otherwise noted, these distributions are parameterized in terms of their mean. The support of Poisson, geometric, and negative binomial distributions are *shifted*, and defined only on positive integers. This less common form of the negative binomial distribution is detailed in the [S1 Appendix](#).

Our main aim here is to characterize, or approximate, forecast distributions. Typically, these are conditional distributions of the form  $\mathbb{P}\{Y_{n+1} \mid Y_{1:n} = y_{1:n}\}$ , or  $\mathbb{P}\{Y_{n+1:n+k} \mid Y_{1:n} = y_{1:n}\}$ . We will let  $\hat{Y}_n$  refer to an estimator available at time  $n$ . Often, it will be an estimator of the one-step-ahead conditional mean  $\mathbb{E}[Y_{n+1} \mid Y_{1:n} = y_{1:n}]$ .

#### 3.2 Croston’s method

Here, we briefly elaborate on Croston’s method and some of its variants, setting up notation for discussions to follow.

In his paper [1], Croston highlighted an important drawback of SES in IDF. SES forecasts, by placing the highest weight on the most recent observation, would lead to the highest forecasts just after a demand observation, and the lowest just before. Heuristically, Croston proposed to separately run SES on interdemand times and positive demand sizes. Concretely, he

set,

$$\hat{M}_{i+1} = \alpha M_{i+1} + (1 - \alpha)\hat{M}_i, \tag{1a}$$

$$\hat{Q}_{i+1} = \alpha Q_{i+1} + (1 - \alpha)\hat{Q}_i. \tag{1b}$$

That is, he proposed to use the EWMA of each series as the forecast estimate. We will denote this recursive computation  $\hat{Q}_{i+1} = \text{EWMA}_\alpha(Q_{1:i+1})$ . Moreover, Croston also discussed possible models for intermittent demands, setting

$$M_i \sim \mathcal{N}(\mu, \sigma^2) \text{ i.i.d.}, \tag{2a}$$

$$Q_i \sim \mathcal{G}(1/\pi) \text{ i.i.d.} \tag{2b}$$

The paper also considers  $\{M_i\} \sim \text{ARIMA}(0, 1, 1)$  with Gaussian innovations.

The discrepancy between the model of (2) and forecasts (1) has been a common issue in the IDF literature that followed. Assuming the model in (2), unbiased forecasts of  $M_{i+1}$ ,  $Q_{i+1}$  are possible simply by setting them to the averages of previous values. Indeed, this is “implied” by the i.i.d. assumption. Instead, however, the forecast heuristic (1) has been praised for its ability to capture serial correlation and nonstationary behavior in both sizes and intervals. This is clearly at odds with the model assumption.

More formally, the EWMA only functions as an *asymptotically* unbiased estimator of the parameter  $\mu$ . Moreover, as noted by Syntetos and Boylan in a series of papers [22, 24, 25], the EWMA of previous interdemand times only results in a biased estimate of  $\pi$ , due to an oversight of the inversion bias  $\mathbb{E}[1/\hat{Q}_i] \neq 1/\mathbb{E}[\hat{Q}_i]$ . In [25], the authors corrected for this bias via a Taylor series approximation about the mean of  $\mathbb{E}[1/\hat{Q}]$  and gave an approximation to an asymptotically unbiased forecast,

$$\hat{Y}_n = (1 - \alpha/2) \frac{\hat{M}_n}{\hat{Q}_n}.$$

Several other variants of forecasts (1) have been explored, *e.g.*, using a simple moving average instead of the EWMA [26, 29].

Both models suggested by Croston allow trajectories with negative and non-integer values. To alleviate this misspecification issue, models with non-Gaussian likelihood have been proposed, *e.g.*, by parameterizing a distribution of positive support with a mean process obeying an IMA model. Nevertheless, such models—namely, *nonnegative EWMA* models—are “ill-fated” [32] since their trajectories converge to 0 over longer terms. Shenstone and Hyndman [18] use this result to point out that no models exist that yield Croston’s forecast estimates as consistent and unbiased estimators while being immune to the convergence problem outlined here.

Exploring possible ways to account for the model–forecast discrepancy in Croston’s method, [17–19, 36] studied numerous alternative probabilistic models. In this paper, we will consider four variations of these models as baseline methods, which we summarize in Table 1. Let us note that, to the best of our knowledge, only the first three models were previously proposed [19, 36].

### 3.3 Renewal processes in discrete time

Renewal processes constitute a central theme in the theory of stochastic processes and their study plays a more general role in probability theory [68, 69]. Broadly, renewal processes are

**Table 1. Baseline models.**

Model	$Q_i$	$M_i$
Static G-Po [1]	$\mathcal{G}(\mu_q)$ , i.i.d.	$\mathcal{PO}(\mu_m)$ , i.i.d.
Static G-NB [19]	$\mathcal{G}(\mu_q)$ , i.i.d.	$\mathcal{NB}(\mu_m, v_m)$ , i.i.d.
EWMA G-Po [18]	$\mathcal{G}(\hat{Q}_{i-1})$ $\hat{Q}_{i-1} = \text{EWMA}_z(Q_{1:i-1})$	$\mathcal{PO}(\hat{M}_{i-1})$ $\hat{M}_{i-1} = \text{EWMA}_z(M_{1:i-1})$
EWMA G-NB	$\mathcal{G}(\hat{Q}_{i-1})$ $\hat{Q}_{i-1} = \text{EWMA}_z(Q_{1:i-1})$	$\mathcal{NB}(\hat{M}_{i-1}, v_m)$ $\hat{M}_{i-1} = \text{EWMA}_z(M_{1:i-1})$

<https://doi.org/10.1371/journal.pone.0259764.t001>

concerned with recurring events in repeated trials, where after each such event, trials start identically from scratch. That is, *interarrival* times between events are independent and identically distributed. For example, in independent spins of a roulette wheel, drawing a certain number is a recurrent event as is observing three consecutive reds.

The example of a roulette wheel may appear unusual, as spins constitute a discrete sequence of outcomes. Renewal processes are mostly introduced as *continuous-time* counting processes (temporal point processes), where inter-arrival times are i.i.d. [68]. Yet, some of the earlier treatments of renewal theory consider the case of recurrent events in discrete time [70, Ch. 13]. We introduce such processes, focusing our attention on the special case where these recurrent events are simple binary outcomes.

**Definition 1.** (Discrete-time renewal process with Bernoulli trials) Let  $\{Z_n\}_n$  define a sequence of (not necessarily independent) binary random variables. Let  $i$  index time steps where  $Z_n = 1$ , as above, defined through the map  $\sigma(i) = \min\{n \mid \sum_{m=1}^n Z_m \geq i\}$ ,  $\sigma(0) = 0$ , and  $Q_i = \sigma(i) - \sigma(i - 1)$ . The sequence  $\{Z_n\}$  defines a discrete-time renewal process (DTRP) if  $Q_i$  are i.i.d.

Our reuse of notation is not coincidental. It should be clear that,

**Remark 1.** The demand arrival process  $\{Y_n > 0\}$  of Static G-Po, and Static G-NB models, as well as Croston’s original models, are DTRPs.

In fact, these models rely on the most basic renewal process—the Bernoulli process—as a demand arrival regime.

Renewal processes, as the name implies, were developed in the context of modeling *failures* in large systems—towards determining when parts of such systems would have to be *renewed*. The underlying intuition is that every time an identical part is *renewed*, the random process with which it fails starts anew.

Thinking in terms of renewal processes paves the way to introducing a class of interpretable, parsimonious and flexible models. It also gives access to tools from renewal theory, such as characterizing multi-step ahead forecast distributions with convolutions of interdemand times, or model fitting via moment-matching methods. As we will see, they will also enable an intuitive connection to temporal point processes, introduced below.

### 3.4 Recurrent neural networks

Recurrent neural networks (RNN) rely on a recurring *hidden state* to model sequential data, and have been widely adopted in time series applications. In contrast to feedforward neural networks, which can be seen as general function approximators, RNNs approximate sequences that (partly) depend on a recurrence relation. See [71, Ch. 10] for an introduction to RNNs. Long short term memory (LSTM) [72] networks, which we use in this paper, define a specific RNN architecture designed to capture long-term interactions.

Concretely, given inputs  $\{x_i\}$ , LSTMs—as in all RNNs—obey the recurrence

$$h_i = \text{LSTM}(h_{i-1}, x_i)$$

where  $h_i \in \mathbb{R}^d$  comprises both the memory (state) and the output of an LSTM cell. Note that in time series applications, inputs often include past time series values. For example, if  $y_i$  denotes the time series of interest,  $x_i = [y_{i-1}, y_{i-2}, \dots]$ . The exact functional form of  $\text{LSTM}(\cdot)$  is composed of several nonlinear projections and compositions of  $h_{i-1}$  and  $x_i$ , that provide mechanisms for representing “long-term memory.” The LSTM network is implemented on virtually all deep learning frameworks such as [73]. For further details on the exact computation and optimization of LSTMs, we refer the reader to the tutorial by Olah [74].

Most previous works in forecasting rely on RNNs to approximate a time series in squared or absolute error. More precisely, the training objective is

$$\begin{aligned} \min_{\Theta} \quad & \|y - \hat{y}\|_2 \\ \text{s.t.} \quad & \hat{y}_i = g(h_i) \quad h_i = \text{LSTM}_{\Theta}(h_{i-1}, x_i), \end{aligned}$$

where  $\Theta$  denotes the set of all parameters of function LSTM, and  $g(\cdot)$  is a suitable projection to the domain of  $y_i$ , such as an affine transformation from  $\mathbb{R}^d$  to  $\mathbb{R}$ . This is the case with previous work on IDF [63, 64]. A more recent approach in neural forecasting considers RNNs not as forecast functions, but to parameterize forecast *distributions* [48], or approximate transition and emission dynamics in an SSM [56]. In this case of *probabilistic* RNNs [75], one assumes

$$y_i \sim \mathcal{D}(\beta) \quad \beta = g_{\beta}(h_i),$$

where  $y_i$  is drawn from a suitable probability distribution denoted by  $\mathcal{D}$  with parameters  $\beta$ . The parameters are then computed from the output of the LSTM, via an appropriate projection to their domains, denoted here by  $g_{\beta}$ . Then, instead of the mean squared error objective, one minimizes the negative log likelihood of the model. In addition to naturally reflecting domain assumptions about  $y_i$ , such as its support or tailedness, this approach naturally leads to forecast distributions. In the sequel, when we use neural networks in the IDF context, we will rely on this approach. Namely, we will let neural networks parameterize conditional interarrival time distributions in self-modulating renewal processes.

### 3.5 Temporal point processes

Temporal point processes (TPP) are probabilistic models for sets of points on the real line. Often, these points represent discrete (instantaneous) *events* in continuous time. Examples of discrete event sets are plenty, such as order arrivals in a financial market, anomalous log entries in a computer system, or earthquakes.

Demand arrivals in a stock control system can be seen as discrete events, and modeled via a TPP. The rationale for this approach is more intuitive for intermittent demands where instances are rare and the main goal is to determine *when* they will occur. Croston [1] observed in his original article that intermittent demand data arose since “updating [occurred] at fixed unit time intervals, which are much shorter than the times between successive demands for the product.” The same approach was later used as a basis in deriving forecast approximations, for example, [16] based their analysis on an underlying Poisson process.

As noted above, DTRPs have better known, continuous time counterparts which constitute a large subclass of TPPs [76, Ch. 4]. As such, discrete-time IDF models we introduced above have natural continuous time counterparts, including those that benefit from recurrent neural networks for modulating interdemand time distributions [77–80].

Then again, discrete-time processes with binary trajectories, such as DTRPs as given in Definition 1, have also been referred to as “point” processes. Determinantal point processes constitute an important example, which define random processes for finite-dimensional binary vectors [81, 82]. The “points” in this case are the vector elements which equal to 1. Although point processes customarily refer to random point collections on a Borel set, this flexible use of nomenclature further justifies the connection between DTRPs and continuous-time models.

Modeling demand directly with a point process is mostly impractical as exact timestamps of demand events are not observed. Moreover, estimation and inference algorithms are computationally intensive: computation times scale asymptotically in the number of events and not the number of demand review periods. However, in increasingly many scenarios such as in e-commerce, the exact demand time is known and can be used directly. In IDF, the number of events is often in the same order as the number of intervals, making computational costs comparable. In this light, we will explore using TPPs on intermittent demand data, and study how directly addressing demand timestamps affects forecast accuracy.

## 4 Models

### 4.1 Discrete-time renewal processes

In previous sections, we discussed how existing IDF models could be cast as renewal processes. We will first follow the obvious extension implied by this generalization, and we will introduce a set of models with more flexible interdemand time distributions. Particularly, we will first consider extending the demand arrival process of *Static* models given above. In this section, we focus on the rationale for doing so.

Previous IDF models have considered the geometric distribution as a model for the times between successive issue points. Doing so is well-justified, it assumes a homogeneous Bernoulli process for demand arrivals, with the probability of observing a demand constant across time and independent of all other intervals. Yet, this is an obvious oversimplification of real-world arrival processes.

A natural generalization of *Static* models is to alter the demand interarrival time distribution. This relaxes the strict assumption imposed by homogeneous Bernoulli arrivals. To frame our analysis, we define the *hazard rate* (cf. the *hazard function* in continuous-time renewal processes),  $h(k)$  for a discrete random variable  $Q$  with c.d.f.  $F$  as

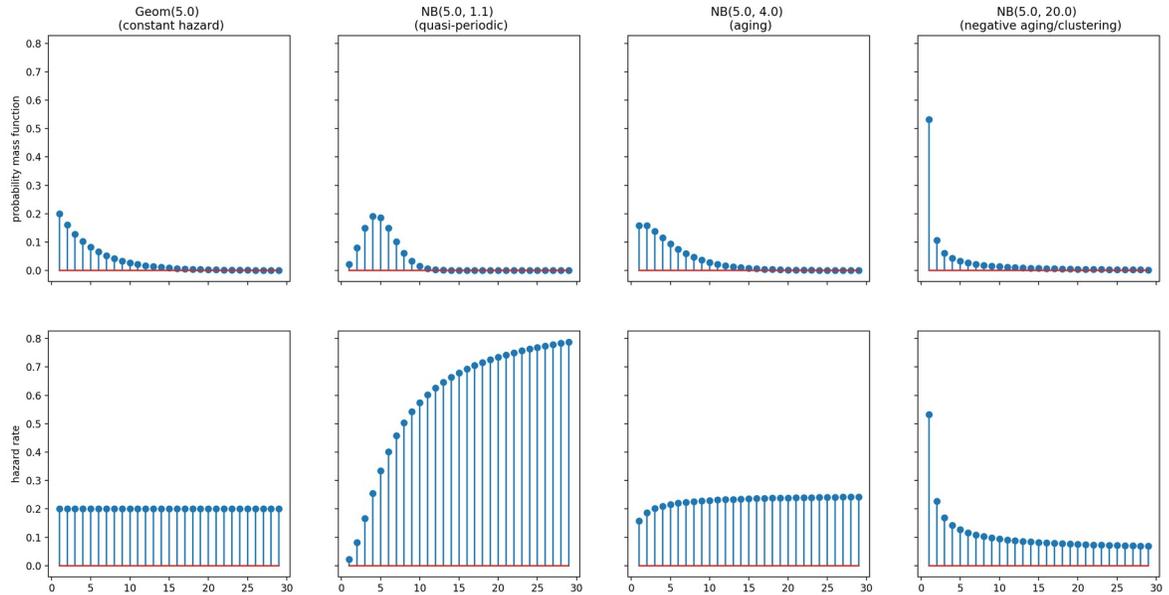
$$h(k) = \mathbb{P}\{Q = k | Q \geq k\} = \mathbb{P}\{Q = k\} / \mathbb{P}\{Q \geq k\} = \frac{\mathbb{P}\{Q = k\}}{1 - F(k-1)}.$$

Letting  $Q$  denote interarrival times, the interpretation of the hazard rate is intuitive. It stands for the probability of a new demand occurring at time  $k$ , given that it hasn't occurred for  $k-1$  intervals since the previous issue point.

It is not hard to show that when  $Q$  admits a geometric distribution, the hazard rate is constant and consequently *independent* of the past. This property of the geometric distribution has been referred to as the *memoryless* property.

If  $h(k)$  is monotonically increasing, the function determines a profile for *aging*. This is often the case in machine parts, where the probability of a part being renewed increases with the time in service.  $h(k)$  can also be decreasing, or imply *negative aging*. In such demand data, issue points are typically *clustered* in time. A flexible distribution for  $Q$  can also be used to capture quasi-periodic demand, e.g., with a distribution for  $Q$  that concentrates around a set period  $\mu$ .

We replace the interarrival time distribution in *Static* models with the negative binomial distribution. The negative binomial can flexibly capture aging, clustering, and periodic effects



**Fig 1. Comparing probability mass functions and hazard rates of geometric and negative binomial distributions.**

<https://doi.org/10.1371/journal.pone.0259764.g001>

discussed above, as illustrated in Fig 1. Indeed, the geometric distribution is a special case of the negative binomial, confirming our intuition that it determines, in a sense, the most basic renewal process possible. We give further details on the negative binomial distribution in the S1 Appendix.

Building on this insight, we introduce the *Static NB-Po* and *Static NB-NB* models, merely by replacing the geometric interarrival time assumption with the negative binomial distribution. As in the rest of this work, we will rely on maximum likelihood parameter estimation and the parametric bootstrap (forward sampling) for forecast estimates of these models. We should note, however, that *Static NB-Po* and *Static NB-NB* models yield closed form forecast estimates, which we explore in the S1 Appendix.

### 4.2 Self-modulating discrete-time renewal processes

The renewal process construction highlights a potentially unrealistic assumption, that interarrival times are independent, while they have been shown to exhibit autocorrelation [15]. Many prior models have either attempted to craft a well-defined probabilistic model for nonstationary and autocorrelated demand arrival regimes, or to encode these relationships via SES-like heuristics [18, 36]. Two examples, *EWMA G-Po* and *EWMA G-NB* were given above.

In order to make a similar extension to these models, we first frame them as instances of *self-modulating* DTRPs. Such models, where the interarrival time distribution is determined by a conditional rate function have been widely explored in the TPP literature [76]. Their key characteristic is that interarrival times now obey a *conditional* distribution, the parameters of which are easily computed. More formally, letting  $Q_i$  denote interarrival times as in Definition 1,

**Definition 2.** A sequence of  $Q_i$  defines a self-modulating DTRP if the sequence obeys an identically defined conditional distribution  $\mathbb{P}\{Q_i = k | \mathcal{H}_i\}$  for all  $i, k$ , where  $\mathcal{H}_i$  denotes the history, of events up to interval  $i$ .

EWMA models are a good example. For all  $i$ , the conditional distribution of  $Q_i$  is defined as a random variable whose mean is the EWMA of previous realizations. In other words, the

history  $\mathcal{H}_i$  completely determines the distribution of  $Q_i$ . A natural next step is to follow the arguments of the previous section, and alter the conditional interarrival time distributions with the negative binomial distribution. We will explore *EWMA NB-Po* and *EWMA NB-NB* models as extensions of their counterparts with geometric interarrival times.

Note that the *convergence problem* that plagues such models is not mitigated by the change in interarrival time distributions. For our model specification, the problem takes a slightly different form. That is, over longer forecast horizons, EWMA model forecasts converge to a forecast of “all ones,” introducing a slightly different bias than the one shown in [32]. We give further details of this issue in [S1 Appendix](#). Finally, note that self-modulating DTRP models can be seen as instances of *Markov renewal processes* [83].

### 4.3 RNN-based discrete-time renewal processes

DTRPs yield an intuitive and simple way to extend IDF models. Yet, they suffer from two main limitations that hinder a realistic understanding of intermittent demand series. First, the conditional distribution of demand sizes and intervals rely only on the exponentially weighted moving average. However, the recurrence relation that determines the conditional mean of sizes and intervals may take more complicated functional forms. The second limitation arises since size and interval processes are assumed independent. One could reasonably expect that longer interdemand intervals result from higher demands—*e.g.*, when a customer *frontloads* inventory—an effect not captured by this independence assumption. Finally, other features or covariates that are consequential for accurate forecasts may be available with each issue point, such as dummy variables for discounts or holidays.

To alleviate the first issue, we observe that under mild conditions any recursive form for parameterizing  $\hat{Q}_i$  would yield a well-defined self-modulating DTRP. An RNN is no exception. Recall from the *EWMA* model definition that

$$\hat{Q}_i = \text{EWMA}_\alpha(Q_{1:i}) = f(\hat{Q}_{i-1}, Q_i) = \alpha Q_i + (1 - \alpha)\hat{Q}_{i-1}.$$

An LSTM, admits a similar recursion, where instead of  $\hat{Q}_i$ , we introduce the *hidden state*  $\mathbf{h}_i \in \mathbb{R}^H$ . We define *RNN-based* models via the following alternative form,

$$\hat{Q}_i = g_w(\mathbf{h}_i), \tag{3}$$

$$\mathbf{h}_i = \text{LSTM}_\Theta(\mathbf{h}_{i-1}, Q_i), \tag{4}$$

where  $g_w(\mathbf{h}) = 1 + \log(1 + \exp(\mathbf{w}^\top \mathbf{h} + w_0))$  is a projection from the hidden state to the domain of  $\hat{Q}_i$ , parameterized by  $\mathbf{w} \in \mathbb{R}^H$ ,  $w_0 \in \mathbb{R}$ . By  $\text{LSTM}_\Theta(\cdot)$ , as given above, we refer to the long short-term memory recurrent neural network function parameterized by weights collected in  $\Theta$ .

Secondly, LSTM inputs can be extended to include both the previous interdemand time, and the previous demand size, at no significant added cost for estimation or inference. Indeed, this is the approach taken by previous IDF models with neural networks [60, 61]. Finally, the LSTM inputs may include any additional covariates that are available, in contrast to other IDF methods that propose no clear way to accommodate them.

Though conceptually similar, EWMA and LSTM recurrences have a fundamental difference. The LSTM network has a potentially high-dimensional parameter vector  $\Theta$  that has to be fitted. Moreover, neural networks are known to be “data-hungry,” requiring a large data set to fit as opposed to the parsimonious representation of temporal continuity embedded in exponential smoothing models. In order to alleviate this issue, when multiple intermittent demand

**Table 2. Discrete time models.**

Model	$Q_i$	$M_i$
Static NB-Po	$\mathcal{NB}(\mu_q, v_q)$ , i.i.d.	$\mathcal{PO}(\mu_m)$ , i.i.d.
Static NB-NB	$\mathcal{NB}(\mu_q, v_q)$ , i.i.d.	$\mathcal{NB}(\mu_m, v_m)$ , i.i.d.
EWMA NB-Po	$\mathcal{NB}(\hat{Q}_{i-1}, v_q)$ $\hat{Q}_{i-1} = \text{EWMA}_\alpha(Q_{1:i-1})$	$\mathcal{PO}(\hat{M}_{i-1})$ $\hat{M}_{i-1} = \text{EWMA}_\alpha(M_{1:i-1})$
EWMA NB-NB	$\mathcal{NB}(\hat{Q}_{i-1}, v_q)$ $\hat{Q}_{i-1} = \text{EWMA}_\alpha(Q_{1:i-1})$	$\mathcal{NB}(\hat{M}_{i-1}, v_m)$ $\hat{M}_{i-1} = \text{EWMA}_\alpha(M_{1:i-1})$
RNN NB-Po	$\mathcal{NB}(\hat{Q}_{i-1}, v_q)$ $\hat{Q}_{i-1} = g(\text{LSTM}_\theta(\mathbf{h}_{i-1}, Q_{i-1}, M_{i-1}))$	$\mathcal{PO}(\hat{M}_{i-1})$ $\hat{M}_{i-1} = g(\text{LSTM}_\theta(\mathbf{h}_{i-1}, Q_{i-1}, M_{i-1}))$
RNN NB-NB	$\mathcal{NB}(\hat{Q}_{i-1}, v_q)$ $\hat{Q}_{i-1} = g(\text{LSTM}_\theta(\mathbf{h}_{i-1}, Q_{i-1}, M_{i-1}))$	$\mathcal{NB}(\hat{M}_{i-1}, v_m)$ $\hat{M}_{i-1} = g(\text{LSTM}_\theta(\mathbf{h}_{i-1}, Q_{i-1}, M_{i-1}))$

<https://doi.org/10.1371/journal.pone.0259764.t002>

series are available, we share the parameters  $\Theta$  across different samples. In machine learning terminology, the LSTM representation learned is a *global* model, shared across different items (SKUs). See, e.g., [47, 59, 84].

We give a summary of the discrete-time renewal process models introduced so far in Table 2.

### 4.4 Continuous time renewal processes

For our last set of models, we explore a connection between the models of the previous section and temporal point processes.

When granular timestamps of individual demand events are available, temporal point processes can be used directly as a model for intermittent demand. A flexible class of continuous-time renewal processes result from the use of RNNs similar to their functions above. To make the connection, note that continuous-time renewal processes arise as a limit case of their discrete time counterparts. For instance, it is well known that the Poisson process—the simplest continuous-time renewal process—arises as a limit case of the Bernoulli process. See, e.g., [85, Ch. 1]. Similarly, the geometric distribution of interarrival times leads to an exponential distribution in the continuous case.

Let us introduce processes  $\{Q'_j\}, \{M'_j\}$  as continuous-time interarrival time and demand size processes. The index  $j$  now runs over individual demand events, e.g., purchase orders, and not only positive demand intervals. We keep the (conditional) distributions of  $M'_j$  identical to the discrete-time model as the support of the random variable and its semantics are identical. To address continuous time, we simply change  $Q'_j$  to a random variable with continuous support. We list these models in Table 3. Note that the *Static E-Po* and *Static E-NB* are just homogeneous Poisson processes with positive integer marks.

This approach, combining RNNs with TPP, has been taken in the machine learning literature to construct flexible models [77–80], with connections to full Markov renewal processes explored recently by [86].

## 5 Experiments

### 5.1 Synthetic data

Our work proposes two main ideas. First, we cast existing IDF models as instances of renewal processes, naturally extending previous point forecast models to probabilistic forecasting. In

**Table 3. Continuous time models.**

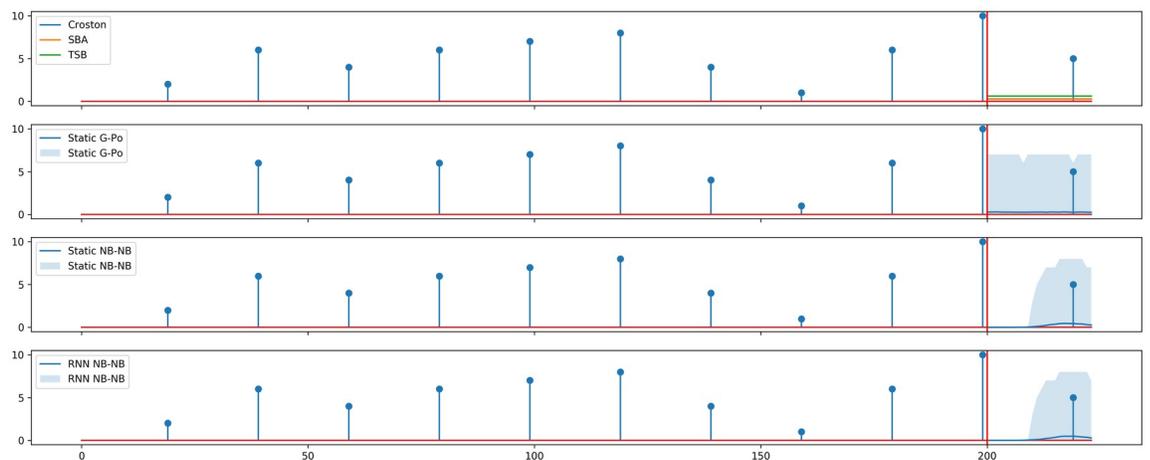
Model	$Q'_i$	$M'_i$
Static E-Po	$\mathcal{E}(\mu_q)$ , i.i.d.	$\mathcal{PO}(\mu_m)$ , i.i.d.
Static E-NB	$\mathcal{E}(\mu_q)$ , i.i.d.	$\mathcal{NB}(\mu_m, v_m)$ , i.i.d.
RNN E-Po	$\mathcal{E}(\hat{Q}_{i-1})$ $\hat{Q}_{i-1} = g(\text{LSTM}_\theta(\mathbf{h}_{i-1}, Q'_{i-1}, M'_{i-1}))$	$\mathcal{PO}(\hat{M}'_{i-1})$ $\hat{M}'_{i-1} = g(\text{LSTM}_\theta(\mathbf{h}_{i-1}, Q'_{i-1}, M'_{i-1}))$
RNN E-NB	$\mathcal{E}(\hat{Q}_{i-1})$ $\hat{Q}_{i-1} = g(\text{LSTM}_\theta(\mathbf{h}_{i-1}, Q'_{i-1}, M'_{i-1}))$	$\mathcal{NB}(\hat{M}'_{i-1})$ $\hat{M}'_{i-1} = g(\text{LSTM}_\theta(\mathbf{h}_{i-1}, Q'_{i-1}, M'_{i-1}))$

<https://doi.org/10.1371/journal.pone.0259764.t003>

contrast to point forecast methods such as Croston’s estimate [1] and the Teunter-Syntetos-Babai (TSB) [30] method, renewal processes are able to produce consistent probabilistic forecasts. Second, by modeling interdemand times with more flexible distributions, we are able to better capture common intermittent demand patterns such as periodicity and aging.

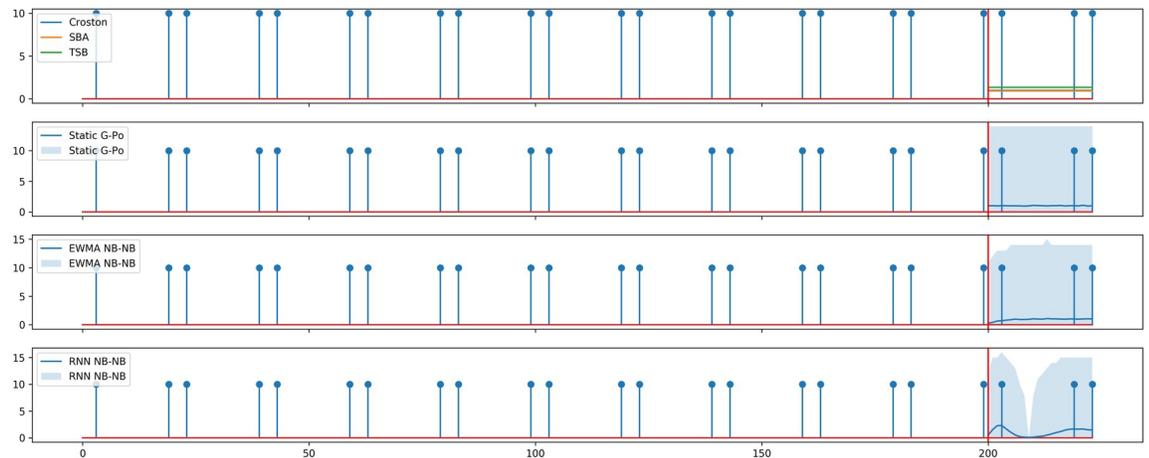
In order to highlight the advantage brought by using flexible renewal processes, we simulate perfectly periodic hourly intermittent demand data of 10 weeks, fixing the interdemand period at 20 hours, and drawing demand sizes i.i.d. from a Poisson distribution with a mean of 5. We illustrate forecasts from previous point forecast methods, Croston [1], Syntetos-Boylan approximation (SBA) [22], and TSB [30] in the top-most panel in Fig 2. The vertical red line marks the beginning of the forecast horizon, and different colored lines in the forecast range refer to point forecasts of different methods, obtained using the `tsintermittent` library [87]. The horizontal red line denotes the x-axis. The next three panels in Fig 2 illustrate forecasts from the Static G-Po [1] model as well as the newly introduced Static NB-NB and RNN NB-NB models respectively. Shaded areas represent 1%–99% forecast intervals obtained by probabilistic forecasting methods via forward sampling—or, parametric bootstrapping. The dark blue line in each graph marks the mean of sampled trajectories. The RNN used is a single-hidden layer LSTM network with five hidden units.

As expected, Static G-Po [1] produces constant forecasts similar to point forecasting methods. On the other hand, by using flexible interdemand time distributions, Static NB-NB and RNN NB-NB models are able to capture periodic demand behavior. Both renewal processes



**Fig 2. Comparing forecasts on a perfectly periodic, synthetic data set.** The vertical red line denotes the start of the forecast period.

<https://doi.org/10.1371/journal.pone.0259764.g002>



**Fig 3. Comparing the forecasts of different IDF approaches on a synthetic data set with alternating periods.**

<https://doi.org/10.1371/journal.pone.0259764.g003>

are able to represent periodicity both in the forecast—i.e., the conditional expectation of the time series—and the forecast quantiles.

Our second contribution is noting that the exponential moving average in Croston-type methods can be replaced with more flexible function families to capture a wider family of patterns. Above, we posit that recurrent neural networks, used as a building block to replace the exponential moving average, can represent many interesting patterns. In order to illustrate this effect, we sample a single time series with “alternating” periods, and a constant demand size of 10. Interdemand times, in turn, alternate between 4 and 16. In Fig 3, we compare point forecasts, Static G-Po, and two new models meant to represent nonstationarities in interdemand times and demand sizes: EWMA NB-NB, and RNN-based NB-NB.

Even in this simple scenario, forecasts of the RNN-based method visibly outperform those of moving average-based methods, as the former is able to easily capture the alternating pattern.

Finally, we validate our findings via experiments on three synthetic data sets. The *Periodic* and *Alternating* data sets are obtained by repeating the respective simulations above 100 times. The *Random* data set is drawn from the Static G-Po model. In a sense, this data set represents purely random demand arrivals, with a Bernoulli process governing the arrival process with Poisson demand sizes.

We expect that our models do not yield any significant benefits in the *Random* data set. Instead, we expect that renewal processes outperform baseline methods in the *Periodic* data set, and the RNN-based model outperforms others in alternating periods. We report results of our experiments, where we compare three point forecast baseline methods (Croston, SBA, TSB), with the probabilistic baseline (Static G-Po) and newly introduced models in Table 4. Definitions of accuracy metrics used can be found in the S1 Appendix.

Among our results, P50 and P90 Losses can be regarded as accuracy measures of probabilistic forecasts while others measure the efficacy of forecasts. We find that our expectation is reflected in the results, and that the RNN model is able to better represent forecast distributions across different data sets.

## 5.2 Real data

We test our framework on five intermittent demand data sets, comparing extended models to well-known baselines. These include three well-known standard data sets for IDF and the recently held M5 competition data set which is mostly composed of intermittent series. We

Table 4. Experiments with synthetic data sets.

Data Sets	Model	P50 Loss	P90 Loss	RMSE	RMSSE
Random	Croston [1]	N/A	N/A	1.535	1.484
	SBA [22]	N/A	N/A	<b>1.534</b>	<b>1.480</b>
	TSB [30]	N/A	N/A	1.571	1.528
	Static G-Po [1]	0.418	0.792	1.537	1.488
	Static NB-NB	0.418	0.796	1.536	1.493
	RNN NB-NB	0.418	<b>0.752</b>	1.534	1.485
Periodic	Croston [1]	N/A	N/A	1.218	1.587
	SBA [22]	N/A	N/A	1.218	1.586
	TSB [30]	N/A	N/A	1.232	1.610
	Static G-Po [1]	0.256	0.460	1.219	1.589
	Static NB-NB	0.256	0.419	<b>1.181</b>	<b>1.536</b>
	RNN NB-NB	0.256	<b>0.412</b>	1.183	1.536
Alternating	Croston [1]	N/A	N/A	2.339	1.692
	SBA [22]	N/A	N/A	2.340	1.691
	TSB [30]	N/A	N/A	2.361	1.709
	Static G-Po [1]	1.022	1.209	2.339	1.691
	Static NB-NB	1.022	1.192	2.332	1.686
	RNN NB-NB	<b>1.021</b>	<b>1.187</b>	<b>2.331</b>	<b>1.686</b>

<https://doi.org/10.1371/journal.pone.0259764.t004>

also introduce a new data set from the UCI Machine Learning repository [88]. We give further details of these data sets below.

- **Car Parts** consists of monthly parts demand data for a US-based automobile company, previously studied in [19, 36]. It includes 2674 time series. We conduct experiments on 2503 time series which have complete data for 51 months, with at least one demand point in the first 45 months.
- The **Auto** data set consists of monthly demand data for 3000 items, for a period of 24 months. It was previously used by [6], and originates from [25]. We report results on 1227 time series that are “intermittent” or “lumpy” based on the Syntetos-Boylan classification (SBC) [25], using the implementation in `tsintermittent` [87].
- The **RAF** data set consists of aerospace parts demand data taken from the Royal Air Force, used previously by [3]. The data set includes 84 months of data for 5000 parts.
- The **UCI** data set includes demand records of retail items of 4070 individual items between 01/12/2010 and 09/12/2011 of a UK-based online retailer [89]. Most notably, the data set is a transaction record of individual purchases, with exact date and timestamps. Among these, we aggregate purchase records of the last three months to daily time series, each corresponding to a single product. We retain time series with demand in the first 30 days, and with less than 90 purchase records in the entire timespan. All 1381 such time series are classified as lumpy demand per SBC.
- The **M5** data set is taken from the recently launched M5 competition [67]. The data set includes daily unit sales per product and store in Walmart, over the course of over 5 years. It contains 30490 time series, with a maximum of 1913 days of data in the longest example. 23401 of these series are classified as intermittent while 5953 are lumpy according to SBC.

Table 5. Summary statistics of data sets.

	Car Parts	Auto	RAF	UCI	M5
Number of Time Series— $M$	2503	1227	5000	1381	30490
Time Series Length— $N$	51	24	84	90	1913
Demand Size—Mean	2.02	3.79	14.19	5.21	3.54
Demand Size— $CV^2$	0.86	4.42	11.89	3.28	3.08
Mean Interdemand Time— $p$	3.41	1.50	8.75	4.89	3.13
$CV^2$ of Interdemand Time	1.93	0.22	0.67	1.72	74.3
Total Number of Issue Points	32,093	56,133	42,695	23,416	18,549,855
Mean Number of Issue Points	12.82	18.71	8.54	16.96	608.4
Number of “Intermittent” Items (SBC)	2087	942	2597	0	23041
Number of “Lumpy” Items (SBC)	412	285	2403	1381	5953

<https://doi.org/10.1371/journal.pone.0259764.t005>

Summary statistics of data sets are given in Table 5, where  $CV^2$  denotes the squared coefficient of variation across all demand instances in the data set.

We take the last 6 demand periods of each time series as a held out sample, and train models on the preceding periods. Results are reported on the forecast accuracy of the last 6 periods, *i. e.*, with a 6-step-ahead forecast. For the M5 competition, both the held-out sample and the forecast horizon are set to 28 days, in accordance with competition rules [67].

We implement all models on Apache MXNet [73], and rely on the Gluon Adam optimizer for parameter fitting wherever necessary [90]. We set the learning rate to 0.1. We run experiments on the Amazon SageMaker platform, cf. [91, 92]. For Croston-type models, we set  $\alpha = 0.1$ . Our implementation of deep renewal processes is available open-source within GluonTS [58]. See, <https://github.com/awsmlabs/gluon-ts/tree/master/src/gluonts/model/renewal>.

Notably, due to the lack of sufficiently large data for cross-validation, we perform no hyperparameter optimization (HPO) for the neural network architecture, regularization, or optimization parameters on the first four data sets. For RNN-based models, we use a single hidden layer of 20 units in the LSTM. For regularization, we rely on *weight decay*, setting the decay rate to 0.01.

On the M5 data set, where sufficiently many examples for cross-validation are available, we perform hyperparameter optimization with Amazon SageMaker’s hyperparameter tuner, which implements Bayesian optimization for searching the hyperparameter space for the minimum value of the loss metric [93]. Here, we heuristically set the loss to be the simple average of the four forecast accuracy metrics reported. We reserve 20% of the in-sample training data as the validation set, and simultaneously tune the learning rate, number of layers, number of hidden units, and number of training epochs. The winning network configuration has 3 hidden layers with 30 hidden units each, which is trained in 750 epochs with a learning rate of 0.01. In our results, we report outcomes for the default hyperparameters given above, and separately for the optimized versions denoting these as “+ HPO.”

We evaluate performance of both forecast distributions and point forecasts. As previously noted, forecast accuracy metrics are especially elusive in the case of IDF. We report results on a variety of metrics which have been used in both IDF and in forecasting in general. The precise definitions of these metrics are given in the S1 Appendix. We repeat each experiment three times, and report the mean and standard deviation of metrics. For probabilistic models, we evaluate forecasts based on 250 trajectories sampled from the fitted model. For point forecast comparisons, we compare true values with means of sampled trajectories for simplicity of experimental design. Choosing the mean as a statistic of the probabilistic forecasts constitutes

a lower bound on the accuracy as choosing different statistics in accordance with the accuracy metric could yield improved results still (see e.g., metric [94] for a discussion). Note that we are interested in evaluating different, distinct methods for addressing the IDF for which we designed our experiment. We remark however that naturally, any of the methods presented in our results can be enhanced with standard techniques such as model selection heuristics (e.g., depending on data set statistics such as intermittency patterns, choose a specific model instantiation) or ensembling techniques which are highly successful in forecasting practice (e.g., [9, 52, 95]).

Discrete-time model results are reported in Tables 6 and 7. *All zeros* refers to accuracy obtained by a naive baseline—a point forecast of 0s across the forecast horizon. Across all metrics, lower numbers indicate better predictive accuracy.

Our results vary across model families and data sets, however some key themes emerge. We observe that in four of five data sets, RNN-based models lead to forecast distributions—as measured with P50 and P90 losses—that are on par with or more accurate than their non-neural counterparts. The exception is the Auto data set, which has a total time series length of 18 steps in sample, barely enough to include any interesting temporal patterns. While these findings support our hypothesis that RNNs can recover more complex temporal patterns, we believe their true potential is offset by the inability to perform hyperparameter optimization. In a real-world scenario, model architecture, training and regularization hyperparameters would be tuned for each specific data set and target forecast horizon. Here, by keeping hyperparameters constant, we report the “bare minimum” of what RNN-based models can achieve in IDF. Indeed, on the M5 data set where hyperparameters were optimized, we find that RNNs significantly outperform all other model-based methods in both probabilistic and point forecast performance, and are slightly better than point forecast methods as measured by RMSE and RMSSE.

Without the use of RNNs, flexible interdemand times alone improve both probabilistic and point forecasts in the first three data sets. This confirms our intuition that improvements by more flexible renewal-type models are data set dependent. The Car Parts data set, which has the lowest variation in demand sizes but a high variation in interdemand intervals yields some evidence in favor of using DTRPs, matching our intuition for where such models would be useful. However, in the larger UCI and M5 data sets, modeling with flexible interdemand sizes alone leads to a deterioration in forecast performance. Finally, we find scant evidence that flexible negative binomial demand sizes improve forecast accuracy.

Our results so far suggest that carefully tuned RNN-based forecasters, built on the DTRP formalism, yield slightly more accurate forecasts. However, we find no strong evidence that DTRPs alone generally outperform baseline methods. Tables 6 and 7 do not paint a clear picture of which modeling direction is more promising. Note, however, that this is in line with other empirical studies on forecasting methods, e.g., [58], where no overall dominant model is found for forecasting—unlike other areas of machine learning such as natural language processing, where dominant models have emerged (e.g., [96]). In order to quantify the improvement brought by our two main modeling ideas, we calculate the significance of improvements brought by model families across different data sets. For this, we compute the ratio of losses for each model to the loss of the Croston (Static G-Po) model, and report averages of these ratios across data sets. In Table 8, we report these average ratios for model families, and the significance level of a one-sample one-sided *t*-test under the null hypothesis that the ratio is equal to or greater than one—*i.e.*, the model brings no improvement.

Here, we find that flexible demand size distributions, proposed in [19] do not generally result in improved forecasts. Moreover, renewal processes—or flexible interdemand sizes alone—do not appear to yield improvements across data sets. We also do not find that

Table 6. Experiment results for discrete-time models with Car Parts, Auto, and RAF data sets.

Data Set	Model	P50 Loss	P90 Loss	RMSE	RMSSE
Car Parts	All Zeros	N/A	N/A	1.513	1.406
	Croston [1]	N/A	N/A	1.372	1.307
	SBA [22]	N/A	N/A	1.368	1.301
	TSB [30]	N/A	N/A	1.362	1.278
	Static G-Po [1]	0.750 ± 0.002	0.638 ± 0.000	1.410 ± 0.001	1.345 ± 0.000
	Static G-NB [19]	0.718 ± 0.001	0.713 ± 0.001	1.434 ± 0.000	1.366 ± 0.001
	EWMA G-Po [18]	0.716 ± 0.002	0.791 ± 0.004	1.438 ± 0.001	1.363 ± 0.000
	EWMA G-NB	0.711 ± 0.001	0.800 ± 0.000	1.438 ± 0.000	1.363 ± 0.001
	Static NB-Po	0.663 ± 0.001	0.553 ± 0.002	1.343 ± 0.002	1.257 ± 0.001
	Static NB-NB	0.663 ± 0.000	0.620 ± 0.003	1.365 ± 0.001	1.290 ± 0.001
	EWMA NB-Po	0.720 ± 0.001	0.751 ± 0.005	1.435 ± 0.002	1.358 ± 0.001
	EWMA NB-NB	0.708 ± 0.001	0.764 ± 0.002	1.436 ± 0.002	1.358 ± 0.001
	RNN G-Po	0.396 ± 0.005	<b>0.447 ± 0.005</b>	<b>1.062 ± 0.004</b>	<b>0.996 ± 0.006</b>
	RNN G-NB	0.401 ± 0.006	0.452 ± 0.009	1.074 ± 0.005	<b>1.015 ± 0.008</b>
	RNN NB-Po	<b>0.388 ± 0.002</b>	0.458 ± 0.003	<b>1.068 ± 0.001</b>	1.021 ± 0.008
RNN NB-NB	<b>0.385 ± 0.001</b>	<b>0.448 ± 0.004</b>	<b>1.068 ± 0.004</b>	1.026 ± 0.027	
Auto	All Zeros	N/A	N/A	9.618	1.820
	Croston [1]	N/A	N/A	7.881	1.477
	SBA [22]	N/A	N/A	7.868	<b>1.468</b>
	TSB [30]	N/A	N/A	8.000	1.501
	Static G-Po [1]	2.383 ± 0.002	<b>1.527 ± 0.001</b>	<b>7.842 ± 0.004</b>	1.479 ± 0.001
	Static G-NB [19]	<b>2.121 ± 0.000</b>	2.277 ± 0.004	8.439 ± 0.003	1.828 ± 0.002
	EWMA G-Po [18]	2.503 ± 0.008	1.574 ± 0.002	8.144 ± 0.010	1.525 ± 0.000
	EWMA G-NB	2.330 ± 0.002	1.605 ± 0.001	8.134 ± 0.005	1.526 ± 0.000
	Static NB-Po	2.335 ± 0.003	<b>1.521 ± 0.001</b>	<b>7.809 ± 0.023</b>	<b>1.452 ± 0.001</b>
	Static NB-NB	<b>2.109 ± 0.002</b>	2.271 ± 0.006	8.465 ± 0.009	1.809 ± 0.000
	EWMA NB-Po	2.493 ± 0.001	1.575 ± 0.001	8.156 ± 0.016	1.524 ± 0.000
	EWMA NB-NB	2.333 ± 0.003	1.606 ± 0.001	8.156 ± 0.008	1.526 ± 0.000
	RNN G-Po	2.395 ± 0.020	1.649 ± 0.010	8.725 ± 0.016	1.492 ± 0.008
	RNN G-NB	2.158 ± 0.028	1.834 ± 0.014	8.814 ± 0.053	1.630 ± 0.031
	RNN NB-Po	2.411 ± 0.036	1.649 ± 0.005	8.634 ± 0.018	1.497 ± 0.016
RNN NB-NB	2.119 ± 0.009	1.892 ± 0.004	8.951 ± 0.015	1.605 ± 0.014	
RAF	All Zeros	N/A	N/A	16.786	<b>1.423</b>
	Croston [1]	N/A	N/A	16.545	1.777
	SBA [22]	N/A	N/A	<b>16.543</b>	1.770
	TSB [30]	N/A	N/A	17.731	<b>1.710</b>
	Static G-Po [1]	1.312 ± 0.000	2.878 ± 0.008	16.557 ± 0.019	1.818 ± 0.000
	Static G-NB [19]	1.287 ± 0.000	<b>2.354 ± 0.001</b>	<b>16.541 ± 0.018</b>	1.922 ± 0.001
	EWMA G-Po [18]	1.339 ± 0.000	3.094 ± 0.006	16.657 ± 0.015	1.915 ± 0.000
	EWMA G-NB	1.320 ± 0.000	2.782 ± 0.005	16.663 ± 0.007	1.925 ± 0.001
	Static NB-Po	1.324 ± 0.000	3.299 ± 0.002	16.607 ± 0.015	1.956 ± 0.000
	Static NB-NB	1.285 ± 0.000	2.448 ± 0.000	16.596 ± 0.021	2.099 ± 0.001
	EWMA NB-Po	1.299 ± 0.004	3.118 ± 0.003	16.666 ± 0.013	1.912 ± 0.001
	EWMA NB-NB	1.291 ± 0.002	2.786 ± 0.006	16.668 ± 0.023	1.921 ± 0.001
	RNN G-Po	<b>1.277 ± 0.001</b>	2.592 ± 0.021	16.625 ± 0.007	1.825 ± 0.012
	RNN G-NB	<b>1.278 ± 0.001</b>	<b>2.407 ± 0.023</b>	16.630 ± 0.006	1.763 ± 0.002
	RNN NB-Po	<b>1.278 ± 0.002</b>	2.893 ± 0.029	16.595 ± 0.010	1.964 ± 0.018
RNN NB-NB	<b>1.278 ± 0.002</b>	2.499 ± 0.005	16.614 ± 0.004	1.837 ± 0.006	

<https://doi.org/10.1371/journal.pone.0259764.t006>

Table 7. Experiment results for discrete-time models with UCI and M5 sets.

Data Set	Model	P50 Loss	P90 Loss	RMSE	RMSSE
UCI	All Zeros	N/A	N/A	7.661	<b>3.413</b>
	Croston [1]	N/A	N/A	7.456	3.466
	SBA [22]	N/A	N/A	7.460	3.462
	TSB [30]	N/A	N/A	7.430	<b>3.409</b>
	Static G-Po [1]	2.059 ± 0.000	3.043 ± 0.006	7.488 ± 0.001	3.514 ± 0.002
	Static G-NB [19]	2.054 ± 0.000	3.041 ± 0.003	7.458 ± 0.001	3.467 ± 0.002
	EWMA G-Po [18]	2.054 ± 0.001	3.079 ± 0.005	7.490 ± 0.001	3.548 ± 0.002
	EWMA G-NB	<b>2.043 ± 0.001</b>	3.126 ± 0.004	7.494 ± 0.001	3.555 ± 0.000
	Static NB-Po	2.059 ± 0.000	3.121 ± 0.004	7.526 ± 0.002	3.585 ± 0.001
	Static NB-NB	2.059 ± 0.000	3.249 ± 0.001	7.504 ± 0.001	3.596 ± 0.003
	EWMA NB-Po	2.120 ± 0.002	3.068 ± 0.004	7.501 ± 0.002	3.568 ± 0.001
	EWMA NB-NB	2.061 ± 0.000	3.113 ± 0.006	7.503 ± 0.001	3.574 ± 0.001
	RNN G-Po	2.071 ± 0.007	<b>2.906 ± 0.004</b>	<b>7.393 ± 0.004</b>	3.562 ± 0.010
	RNN G-NB	<b>2.039 ± 0.001</b>	<b>2.870 ± 0.004</b>	<b>7.368 ± 0.004</b>	3.566 ± 0.012
	RNN NB-Po	2.086 ± 0.005	2.937 ± 0.005	7.423 ± 0.004	3.573 ± 0.001
RNN NB-NB	2.062 ± 0.003	2.923 ± 0.004	7.417 ± 0.004	3.539 ± 0.018	
M5	All Zeros	N/A	N/A	3.852	1.730
	Croston [1]	N/A	N/A	2.299	1.370
	SBA [22]	N/A	N/A	2.280	1.364
	TSB [30]	N/A	N/A	2.213	<b>1.323</b>
	Static G-Po [1]	1.093 ± 0.000	0.730 ± 0.000	2.453 ± 0.000	1.414 ± 0.000
	Static G-NB [19]	1.250 ± 0.000	1.192 ± 0.001	3.255 ± 0.001	1.520 ± 0.000
	EWMA G-Po [18]	1.031 ± 0.000	0.667 ± 0.000	2.301 ± 0.000	1.377 ± 0.000
	EWMA G-NB	1.015 ± 0.000	0.664 ± 0.000	2.301 ± 0.000	1.378 ± 0.000
	Static NB-Po	1.378 ± 0.000	0.786 ± 0.000	3.040 ± 0.002	1.529 ± 0.000
	Static NB-NB	1.384 ± 0.000	1.550 ± 0.001	3.578 ± 0.001	1.728 ± 0.000
	EWMA NB-Po	1.055 ± 0.000	0.663 ± 0.000	2.306 ± 0.001	1.398 ± 0.000
	EWMA NB-NB	1.032 ± 0.000	0.661 ± 0.000	2.308 ± 0.000	1.399 ± 0.000
	RNN G-Po	1.069 ± 0.051	0.674 ± 0.011	2.554 ± 0.112	1.381 ± 0.026
	RNN G-NB	1.062 ± 0.026	0.720 ± 0.016	2.638 ± 0.038	1.369 ± 0.011
	RNN NB-Po	1.080 ± 0.031	0.666 ± 0.004	2.505 ± 0.097	1.393 ± 0.015
	RNN NB-NB	1.022 ± 0.001	0.711 ± 0.008	2.522 ± 0.046	1.395 ± 0.010
	RNN G-Po + HPO	<b>0.970 ± 0.005</b>	0.647 ± 0.007	<b>2.190 ± 0.011</b>	1.338 ± 0.003
	RNN G-NB + HPO	<b>0.968 ± 0.006</b>	0.639 ± 0.001	<b>2.203 ± 0.012</b>	<b>1.328 ± 0.001</b>
	RNN NB-Po + HPO	1.011 ± 0.014	<b>0.638 ± 0.001</b>	2.245 ± 0.065	1.354 ± 0.007
RNN NB-NB + HPO	0.999 ± 0.006	<b>0.638 ± 0.000</b>	2.235 ± 0.013	1.353 ± 0.001	

<https://doi.org/10.1371/journal.pone.0259764.t007>

exponential moving average-based models result in any significant improvement. Comparing Table 8 to dataset results supports our claim that while these three model components may improve individual data set performance, they do not do so in general. However, we can confirm that these approaches improve on the P50 loss by a slight margin of 2 to 5 per cent.

On the other hand, RNN-based models result in significantly more accurate point and probabilistic forecasts compared to simple probabilistic models. RNNs yield an improvement of 10 per cent on P50 loss, and over 5 per cent in P90 loss in general, yielding more accurate probabilistic forecasts. Their performance in squared losses also outperforms the probabilistic baseline. We can confirm from Tables 6 and 7, however, that RNN models do not generally

**Table 8. Ratios of model families to the Static G-Po baseline model.**

Model Feature	P50 Loss	P90 Loss	RMSE	RMSSE
Flexible Demand Size Distribution (*-NB Models)	0.952 (0.000*)	1.083 (—)	1.023 (—)	1.028 (—)
Flexible Interdemand Time Distribution (NB-* Models)	0.966 (0.000*)	1.053 (—)	1.015 (—)	1.020 (—)
EWMA	0.983 (0.000*)	1.043 (—)	1.001 (—)	1.019 (—)
RNN	0.896 (0.000*)	0.945 (0.000*)	0.973 (0.002*)	0.967 (0.000*)
RNN & Flexible Demand Size	0.884 (0.000*)	0.948 (0.002*)	0.976 (0.036*)	0.970 (0.008*)
RNN & Flexible Interdemand Time	0.895 (0.000*)	0.956 (0.004*)	0.974 (0.023*)	0.973 (0.013*)

<https://doi.org/10.1371/journal.pone.0259764.t008>

yield better point forecasts than classical point forecasting methods. We therefore conclude that while RNNs are promising as probabilistic forecasting methods, further work has to be invested in their use before their performance is on par with carefully tuned and debiased point forecasters such as SBA and TSB. Nevertheless, results from the M5 data set are promising, as the margins between tuned RNN models and point forecasters are small, with the former slightly outperforming the latter.

Finally, we test using TPPs, or continuous-time renewal processes, directly for IDF. The UCI data set is the only one where purchase records with exact timestamps are available. In Table 9, we compare Static and RNN-based models with their continuous time counterparts. Specifically, we train continuous-time models on interdemand times of individual purchase orders instead of the aggregated intermittent demand series. We take forward samples from these models, and aggregate the predicted purchase events in the last 6 time periods. We compare these with the forecast accuracy of samples taken from discrete-time models. Our results yield no evidence in favor of continuous-time models, which appear to exhibit a similar performance with their discrete-time analogues.

## 6 Discussion and conclusion

IDF is a uniquely challenging problem; the definition of good forecasts is as elusive as the techniques used to produce them. Most previous works in IDF, starting from Croston's method [1], have focused on point forecast methods with strict assumptions. While probabilistic versions of Croston's method were introduced, these did not propose a general theoretical framework for similarly constructed forecasting methods.

**Table 9. Experiment results for continuous-time models.**

Model	P50 Loss	P90 Loss	RMSE	RMSSE
Static G-Po [1]	2.059	3.043	7.488	3.514
Static G-NB [19]	2.054	3.041	7.458	<b>3.467</b>
RNN G-Po	2.071	2.906	7.393	3.562
RNN G-NB	<b>2.039</b>	<b>2.870</b>	7.368	3.566
Static E-Po (cont.)	2.062	3.039	7.489	3.518
Static E-NB (cont.)	2.055	3.015	7.449	3.477
RNN E-Po (cont.)	2.055	2.937	7.416	3.544
RNN E-NB (cont.)	2.059	2.871	<b>7.359</b>	3.738

<https://doi.org/10.1371/journal.pone.0259764.t009>

In this work, we proposed a general probabilistic framework for building IDF models, by making connections to well-established literatures in renewal theory and probabilistic neural forecasting. We argued that previously proposed model-based and point forecast methods can all be cast as special cases of our framework, and that it allows for significant generalizations of these methods. For example, by extending simple, widely used models with flexible inter-demand time distributions, we are able to recover various common demand arrival patterns such as aging, periodicity, and clustering. RNNs, used as a subcomponent of our framework, are able to capture non-trivial temporal patterns, *e.g.*, alternating periods which we illustrated on synthetic data.

We reported mixed results from a large set of numerical experiments on a variety of well-known intermittent demand data sets. We found that while renewal processes alone can only lead to improved probabilistic forecasts for individual data sets, they do not generally lead to substantial improvements in performance. However, we found some evidence that RNNs, used as a subcomponent in our framework, lead to improvements in probabilistic forecast accuracy. These improvements are promising for the larger data sets, such as the M5 competition which lend themselves particularly well to neural network based techniques and are examples for an important class of practical forecasting problems [97].

Our study also highlights the intimate connection between temporal point processes and intermittent demand models. Neural TPPs (see *e.g.*, [98]), a recent research direction in machine learning, can be seen as continuous-time instances of our framework. These models are especially well-suited to scenarios in which purchase order data are directly available, removing the need to temporally aggregate demand instances, and directly taking advantage of full temporal granularity. While our empirical results on limited data available do not yield substantive evidence for taking this approach, we believe this connection can be pursued further for demand forecasting research.

In the context of probabilistic forecasting, our theoretical framework does not only yield useful forecasting tools but also opens up exciting avenues for further research for which especially in the recent data set from the M5 competition will provide inspiration. To provide two concrete examples: (i) we believe that a better understanding of the intermittency patterns in the M5 data set will allow for an improved, more fine-tuned model selection of the different instantiations of our framework (or additionally, more flexible extensions thereof); and (ii) combination of the probabilistic models provided here with hierarchical forecasting techniques for probabilistic methods (such as [99, 100]) will offer ample opportunity for further work. In more general terms, although we do not find that the individual models categorically outperform classical point forecasting baselines, we believe experimenting with other instances of our framework, *e.g.*, with other statistical distributions, neural network architectures and training strategies, can lead to closing this gap. Finally, we hope that casting intermittent demand problems in the language of renewal theory can lead to new results and simple tools, *e.g.*, for multi-step ahead forecast estimates or generalized point forecast methods.

## Supporting information

### S1 Appendix.

(ZIP)

## Acknowledgments

The authors would like to thank Nikolaos Kourentzes for kindly providing the Auto and RAF data sets.

## Author Contributions

**Conceptualization:** Ali Caner Türkmen, Tim Januschowski, Yuyang Wang, Ali Taylan Cemgil.

**Data curation:** Ali Caner Türkmen.

**Formal analysis:** Ali Caner Türkmen, Tim Januschowski, Yuyang Wang.

**Investigation:** Ali Caner Türkmen, Tim Januschowski.

**Methodology:** Ali Caner Türkmen, Yuyang Wang, Ali Taylan Cemgil.

**Project administration:** Ali Caner Türkmen, Tim Januschowski, Yuyang Wang.

**Software:** Ali Caner Türkmen.

**Supervision:** Tim Januschowski, Yuyang Wang, Ali Taylan Cemgil.

**Validation:** Ali Caner Türkmen.

**Visualization:** Ali Caner Türkmen.

**Writing – original draft:** Ali Caner Türkmen, Yuyang Wang.

**Writing – review & editing:** Ali Caner Türkmen, Tim Januschowski, Yuyang Wang, Ali Taylan Cemgil.

## References

1. Croston JD. Forecasting and stock control for intermittent demands. *Journal of the Operational Research Society*. 1972; 23(3):289–303. <https://doi.org/10.1057/jors.1972.50>
2. Williams TM. Stock control with sporadic and slow-moving demand. *Journal of the Operational Research Society*. 1984; 35(10):939–948. <https://doi.org/10.2307/2582137>
3. Eaves AH, Kingsman BG. Forecasting for the ordering and stock-holding of spare parts. *Journal of the Operational Research Society*. 2004; 55(4):431–437. <https://doi.org/10.1057/palgrave.jors.2601697>
4. Boylan JE, Syntetos AA, Karakostas GC. Classification for forecasting and stock control: a case study. *Journal of the Operational Research Society*. 2008; 59(4):473–481. <https://doi.org/10.1057/palgrave.jors.2602312>
5. Hyndman RJ. Another look at forecast-accuracy metrics for intermittent demand. *Foresight: The International Journal of Applied Forecasting*. 2006; 4(4):43–46.
6. Kourentzes N. On intermittent demand model optimisation and selection. *International Journal of Production Economics*. 2014; 156:180–190. <https://doi.org/10.1016/j.ijpe.2014.06.007>
7. Petropoulos F, Kourentzes N. Forecast combinations for intermittent demand. *Journal of the Operational Research Society*. 2015; 66(6):914–924. <https://doi.org/10.1057/jors.2014.62>
8. Seeger MW, Salinas D, Flunkert V. Bayesian intermittent demand forecasting for large inventories. In: *Advances in Neural Information Processing Systems*; 2016. p. 4646–4654.
9. Böse JH, Flunkert V, Gasthaus J, Januschowski T, Lange D, Salinas D, et al. Probabilistic demand forecasting at scale. *Proceedings of the VLDB Endowment*. 2017; 10(12):1694–1705. <https://doi.org/10.14778/3137765.3137775>
10. Syntetos AA, Boylan JE. Intermittent demand: estimation and statistical properties. In: *Service Parts Management*. Springer; 2011. p. 1–30.
11. Bacchetti A, Sacconi N. Spare parts classification and demand forecasting for stock control: Investigating the gap between research and practice. *Omega*. 2012; 40(6):722–737. <https://doi.org/10.1016/j.omega.2011.06.008>
12. Ward JB. Determining Reorder Points When Demand is Lumpy. *Management Science*. 1978; 24(6):623–632. <https://doi.org/10.1287/mnsc.24.6.623>
13. Schultz CR. Forecasting and Inventory Control for Sporadic Demand Under Periodic Review. *Journal of the Operational Research Society*. 1987; 38(5):453–458. <https://doi.org/10.1057/jors.1987.74>
14. Dunsmuir WTM, Snyder RN. Control of inventories with intermittent demand. *European Journal of Operational Research*. 1989; 40(1):16–21. [https://doi.org/10.1016/0377-2217\(89\)90267-1](https://doi.org/10.1016/0377-2217(89)90267-1)

15. Willemain TR, Smart CN, Shockor JH, DeSautels PA. Forecasting intermittent demand in manufacturing: a comparative evaluation of Croston's method. *International Journal of forecasting*. 1994; 10(4):529–538. [https://doi.org/10.1016/0169-2070\(94\)90021-3](https://doi.org/10.1016/0169-2070(94)90021-3)
16. Johnston FR, Boylan JE. Forecasting for items with intermittent demand. *Journal of the operational research society*. 1996; 47(1):113–121. <https://doi.org/10.1057/jors.1996.10>
17. Snyder R. Forecasting sales of slow and fast moving inventories. *European Journal of Operational Research*. 2002; 140(3):684–699. [https://doi.org/10.1016/S0377-2217\(01\)00231-4](https://doi.org/10.1016/S0377-2217(01)00231-4)
18. Shenstone L, Hyndman RJ. Stochastic models underlying Croston's method for intermittent demand forecasting. *Journal of Forecasting*. 2005; 24(6):389–402. <https://doi.org/10.1002/for.963>
19. Snyder RD, Ord JK, Beaumont A. Forecasting the intermittent demand for slow-moving inventories: A modelling approach. *International Journal of Forecasting*. 2012; 28(2):485–496. <https://doi.org/10.1016/j.ijforecast.2011.03.009>
20. Türkmen AC, Wang Y, Januschowski T. Intermittent Demand Forecasting with Deep Renewal Processes. In: *NeurIPS Workshop on Learning with Temporal Point Processes*; 2019.
21. Rao AV. A Comment on: Forecasting and Stock Control for Intermittent Demands. *Journal of the Operational Research Society*. 1973; 24(4):639–640. <https://doi.org/10.1057/jors.1973.120>
22. Syntetos AA, Boylan JE. On the bias of intermittent demand estimates. *International journal of production economics*. 2001; 71(1-3):457–466. [https://doi.org/10.1016/S0925-5273\(00\)00143-2](https://doi.org/10.1016/S0925-5273(00)00143-2)
23. Hyndman RJ, Athanasopoulos G. *Forecasting: principles and practice*. OTexts; 2018.
24. Syntetos A. *Forecasting of intermittent demand [PhD Thesis]*. Brunel University Uxbridge; 2001.
25. Syntetos AA, Boylan JE. The accuracy of intermittent demand estimates. *International Journal of forecasting*. 2005; 21(2):303–314. <https://doi.org/10.1016/j.ijforecast.2004.10.001>
26. Shale EA, Boylan JE, Johnston FR. Forecasting for intermittent demand: the estimation of an unbiased average. *Journal of the Operational Research Society*. 2006; 57(5):588–592. <https://doi.org/10.1057/palgrave.jors.2602031>
27. Syntetos AA, Boylan JE, Croston JD. On the categorization of demand patterns. *Journal of the Operational Research Society*. 2005; 56(5):495–503. <https://doi.org/10.1057/palgrave.jors.2601841>
28. Kostenko AV, Hyndman RJ. A note on the categorization of demand patterns. *Journal of the Operational Research Society*. 2006; 57(10):1256–1257. <https://doi.org/10.1057/palgrave.jors.2602211>
29. Boylan JE, Syntetos AA. Forecasting for inventory management of service parts. In: *Complex system maintenance handbook*. Springer; 2008. p. 479–506.
30. Teunter RH, Syntetos AA, Babai MZ. Intermittent demand: Linking forecasting to inventory obsolescence. *European Journal of Operational Research*. 2011; 214(3):606–615. <https://doi.org/10.1016/j.ejor.2011.05.018>
31. Babai MZ, Syntetos A, Teunter R. Intermittent demand forecasting: An empirical study on accuracy and the risk of obsolescence. *International Journal of Production Economics*. 2014; 157:212–219. <https://doi.org/10.1016/j.ijpe.2014.08.019>
32. Grunwald GK, Hamza K, Hyndman RJ. Some properties and generalizations of non-negative Bayesian time series models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 1997; 59(3):615–626. <https://doi.org/10.1111/1467-9868.00086>
33. Akram M, Hyndman RJ, Ord JK. Exponential smoothing and non-negative data. *Australian & New Zealand Journal of Statistics*. 2009; 51(4):415–432. <https://doi.org/10.1111/j.1467-842X.2009.00555.x>
34. Dolgui A, Pashkevich A, Pashkevich M. Bayesian approach to modelling of quasi-periodic intermittent demand. *IFAC Proceedings Volumes*. 2005; 38(1):343–348. <https://doi.org/10.3182/20050703-6-CZ-1902.01479>
35. Sandmann W, Bober O. Stochastic Models For Intermittent demands forecasting and stock control. In: *Proceedings of the Vienna Conference on Mathematical Modelling, MATHMOD, Vienna, Austria*; 2009.
36. Hyndman RJ, Koehler AB, Ord JK, Snyder RD. *Forecasting with exponential smoothing: the state space approach*. Springer series in statistics. Berlin: Springer; 2008.
37. Yelland PM. Bayesian forecasting for low-count time series using state-space models: An empirical evaluation for inventory management. *International Journal of Production Economics*. 2009; 118(1):95–103. <https://doi.org/10.1016/j.ijpe.2008.08.027>
38. Yelland PM. Bayesian forecasting of parts demand. *International Journal of Forecasting*. 2010; 26(2):374–396. <https://doi.org/10.1016/j.ijforecast.2009.11.001>
39. Seeger M, Rangapuram S, Wang Y, Salinas D, Gasthaus J, Januschowski T, et al. Approximate Bayesian Inference in Linear State Space Models for Intermittent Demand Forecasting at Scale; 2017.

40. Altay N, Rudisill F, Litteral LA. Adapting Wright's modification of Holt's method to forecasting intermittent demand. *International Journal of Production Economics*. 2008; 111(2):389–408. <https://doi.org/10.1016/j.ijpe.2007.01.009>
41. Gamberini R, Lolli F, Rimini B, Sgarbossa F. Forecasting of sporadic demand patterns with seasonality and trend components: an empirical comparison between Holt-Winters and (S) ARIMA methods. *Mathematical Problems in Engineering*. 2010; 2010. <https://doi.org/10.1155/2010/579010>
42. Svetunkov I, Boylan JE. Multiplicative state-space models for intermittent time series. Lancaster University Management School; 2017.
43. Kourentzes N, Athanasopoulos G. Elucidate Structure in Intermittent Demand Series. Monash University, Department of Econometrics and Business Statistics; 2019.
44. Chatfield C, et al. Neural networks: Forecasting breakthrough or passing fad? *International Journal of Forecasting*. 1993; 9(1):1–3. [https://doi.org/10.1016/0169-2070\(93\)90043-M](https://doi.org/10.1016/0169-2070(93)90043-M)
45. Crone SF, Hibon M, Nikolopoulos K. Advances in forecasting with neural networks? Empirical evidence from the NN3 competition on time series prediction. *International Journal of forecasting*. 2011; 27(3):635–660. <https://doi.org/10.1016/j.ijforecast.2011.04.001>
46. Makridakis S, Spiliotis E, Assimakopoulos V. Statistical and Machine Learning forecasting methods: Concerns and ways forward. *PloS one*. 2018; 13(3):e0194889. <https://doi.org/10.1371/journal.pone.0194889> PMID: 29584784
47. Januschowski T, Gasthaus J, Wang Y, Salinas D, Flunkert V, Bohlke-Schneider M, et al. Criteria for classifying forecasting methods. *International Journal of Forecasting*. 2019.
48. Salinas D, Flunkert V, Gasthaus J, Januschowski T. DeepAR: Probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting*. 2019.
49. Laptev N, Yosinsk J, Li Erran L, Smyl S. Time-series Extreme Event Forecasting with Neural Networks at Uber. In: *ICML Time Series Workshop*; 2017.
50. Wen R, Torkkola K, Narayanaswamy B, Madeka D. A Multi-Horizon Quantile Recurrent Forecaster; 2017.
51. Bandara K, Shi P, Bergmeir C, Hewamalage H, Tran Q, Seaman B. Sales Demand Forecast in E-commerce using a Long Short-Term Memory Neural Network Methodology. *CoRR*. 2019;abs/1901.04028.
52. Makridakis S, Spiliotis E, Assimakopoulos V. The M4 Competition: Results, findings, conclusion and way forward. *International Journal of Forecasting*. 2018; 34(4):802–808. <https://doi.org/10.1016/j.ijforecast.2018.06.001>
53. Smyl S. A hybrid method of exponential smoothing and recurrent neural networks for time series forecasting. *International Journal of Forecasting*. 2020; 36(1):75–85. <https://doi.org/10.1016/j.ijforecast.2019.03.017>
54. Kourentzes N, Barrow DK, Crone SF. Neural network ensemble operators for time series forecasting. *Expert Systems with Applications*. 2014; 41(9):4235–4244. <https://doi.org/10.1016/j.eswa.2013.12.011>
55. Wang Y, Smola A, Maddix DC, Gasthaus J, Foster D, Januschowski T. Deep factors for forecasting. *arXiv preprint arXiv:190512417*. 2019.
56. Rangapuram SS, Seeger MW, Gasthaus J, Stella L, Wang Y, Januschowski T. Deep state space models for time series forecasting. In: *Advances in Neural Information Processing Systems*; 2018. p. 7785–7794.
57. Gasthaus J, Benidis K, Wang Y, Rangapuram SS, Salinas D, Flunkert V, et al. Probabilistic forecasting with spline quantile function rns. In: *The 22nd International Conference on Artificial Intelligence and Statistics*; 2019. p. 1901–1910.
58. Alexandrov A, Benidis K, Bohlke-Schneider M, Flunkert V, Gasthaus J, Januschowski T, et al. GluonTS: Probabilistic time series models in python. *Journal of Machine Learning Research*. 2020.
59. Benidis K, Rangapuram SS, Flunkert V, Wang B, Maddix D, Turkmen C, et al. Neural forecasting: Introduction and literature overview. *arXiv preprint arXiv:200410240*. 2020.
60. Gutierrez RS, Solis AO, Mukhopadhyay S. Lumpy demand forecasting using neural networks. *International Journal of Production Economics*. 2008; 111(2):409–420. <https://doi.org/10.1016/j.ijpe.2007.01.007>
61. Kourentzes N. Intermittent demand forecasts with neural networks. *International Journal of Production Economics*. 2013; 143(1):198–206. <https://doi.org/10.1016/j.ijpe.2013.01.009>
62. Mukhopadhyay S, Solis AO, Gutierrez RS. The accuracy of non-traditional versus traditional methods of forecasting lumpy demand. *Journal of Forecasting*. 2012; 31(8):721–735. <https://doi.org/10.1002/for.1242>

63. Nasiri Pour A, Rostami Tabar B, Rahimzadeh A. A hybrid neural network and traditional approach for forecasting lumpy demand. *Proc World Acad Sei Eng Technol*. 2008; 30:384–389.
64. Lolli F, Gamberini R, Regattieri A, Balugani E, Gatos T, Gucci S. Single-hidden layer neural networks for forecasting intermittent demand. *International Journal of Production Economics*. 2017; 183:116–128. <https://doi.org/10.1016/j.ijpe.2016.10.021>
65. Fu W, Chien CF, Lin ZH. A hybrid forecasting framework with neural network and time-series method for intermittent demand in semiconductor supply chain. In: *IFIP International Conference on Advances in Production Management Systems*. Springer; 2018. p. 65–72.
66. Kim S, Kim H. A new metric of absolute percentage error for intermittent demand forecasts. *International Journal of Forecasting*. 2016; 32(3):669–679. <https://doi.org/10.1016/j.ijforecast.2015.12.003>
67. Makridakis S. *The M5 Competition Competitors' Guide*. University of Nicosia; 2020. Available from: <https://www.kaggle.com/c/m5-forecasting-accuracy/data>.
68. Cox DR. *Renewal theory*. John Wiley and Sons; 1962.
69. Çinlar E. *Introduction to stochastic processes*. Courier Corporation; 2013.
70. Feller W. *An introduction to probability theory and its applications*. John Wiley & Sons; 1957.
71. Goodfellow I, Bengio Y, Courville A. *Deep learning*. MIT press; 2016.
72. Hochreiter S, Schmidhuber J. Long short-term memory. *Neural computation*. 1997; 9(8):1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735> PMID: 9377276
73. Chen T, Li M, Li Y, Lin M, Wang N, Wang M, et al. Mxnet: A flexible and efficient machine learning library for heterogeneous distributed systems. *arXiv preprint arXiv:151201274*. 2015.
74. Olah C. *Understanding lstm networks*; 2015. Available from: <http://colah.github.io/posts/2015-08-Understanding-LSTMs/>.
75. Januschowski T, Gasthaus J, Wang Y, Salinas D, Flunkert V, Bohlke-Schneider M, et al. Criteria for classifying forecasting methods. *International Journal of Forecasting*. 2020; 36(1):167–177. <https://doi.org/10.1016/j.ijforecast.2019.05.008>
76. Daley D, Vere-Jones D. *An Introduction to the Theory of Point Processes: Volume I: Elementary Theory and Methods*; 2003.
77. Du N, Dai H, Trivedi R, Upadhyay U, Gomez-Rodriguez M, Song L. Recurrent marked temporal point processes: Embedding event history to vector. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM; 2016. p. 1555–1564.
78. Mei H, Eisner JM. The neural Hawkes process: A neurally self-modulating multivariate point process. In: *Advances in Neural Information Processing Systems*; 2017. p. 6754–6764.
79. Türkmen AC, Wang Y, Smola AJ. FastPoint: Scalable Deep Point Processes. In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*; 2019.
80. Shchur O, Bilos M, Günnemann S. Intensity-Free Learning of Temporal Point Processes. In: *International Conference on Learning Representations*; 2019.
81. Borodin A. Determinantal point processes. In: *The Oxford Handbook of Random Matrix Theory*.
82. Kulesza A, Taskar B, et al. Determinantal point processes for machine learning. *Foundations and Trends in Machine Learning*. 2012; 5(2–3):123–286. <https://doi.org/10.1561/22000000044>
83. Çinlar E. Markov Renewal Theory: A Survey. *Management Science*. 1975; 21(7):27.
84. Januschowski T, Gasthaus J, Wang Y, Rangapuram SS, Callot L. *Deep Learning for Forecasting: Current Trends and Challenges*. *Foresight: The International Journal of Applied Forecasting*. 2018; 51.
85. Kingman J. *Poisson Processes*. Clarendon Press; 1992.
86. Sharma A, Johnson R, Engert F, Linderman S. Point process latent variable models of larval zebrafish behavior. In: *Advances in Neural Information Processing Systems*; 2018. p. 10919–10930.
87. Kourentzes N, Petropoulos F. *tsintermittent: Intermittent time series forecasting*; 2014. R Package.
88. Asuncion A, Newman D. *UCI machine learning repository*; 2007.
89. Chen D, Sain SL, Guo K. Data mining for the online retail industry: A case study of RFM model-based customer segmentation using data mining. *Journal of Database Marketing & Customer Strategy Management*. 2012; 19(3):197–208. <https://doi.org/10.1057/dbm.2012.17>
90. Kingma DP, Ba J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*. 2014.
91. Januschowski T, Arpin D, Salinas D, Flunkert V, Gasthaus J, Stella L, et al. Now available in amazon sagemaker: Deepar algorithm for more accurate time series forecasting; 2018. *AWS Machine Learning Blog*.

92. Liberty E, Karnin Z, Xiang B, Rouesnel L, Coskun B, Nallapati R, et al. Elastic Machine Learning Algorithms in Amazon SageMaker. In: Proceedings of the 2020 International Conference on Management of Data. SIGMOD'20. New York, NY, USA: Association for Computing Machinery; 2020.
93. Perrone V, Shen H, Zolic A, Shcherbatyi I, Ahmed A, Bansal T, et al. Amazon SageMaker Automatic Model Tuning: Scalable Black-box Optimization; 2020.
94. Kolassa S. Why the “best” point forecast depends on the error or accuracy measure. *International Journal of Forecasting*. 2020; 36(1):208–211. <https://doi.org/10.1016/j.ijforecast.2019.02.017>
95. Petropoulos F, Apiletti D, Assimakopoulos V, Babai MZ, Barrow DK, Taieb SB, et al. Forecasting: theory and practice; 2021.
96. Devlin J, Chang MW, Lee K, Toutanova K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding; 2018.
97. Januschowski T, Kolassa S. A Classification of Business Forecasting Problems. *Foresight: The International Journal of Applied Forecasting*. 2019;(52):36–43.
98. Shchur O, Türkmen AC, Januschowski T, Günemann S. Neural Temporal Point Processes: A Review; 2021.
99. Rangapuram SS, Werner LD, Benidis K, Mercado P, Gasthaus J, Januschowski T. End-to-End Learning of Coherent Probabilistic Forecasts for Hierarchical Time Series. In: Meila M, Zhang T, editors. Proceedings of the 38th International Conference on Machine Learning. vol. 139 of Proceedings of Machine Learning Research. PMLR; 2021. p. 8832–8843. Available from: <https://proceedings.mlr.press/v139/rangapuram21a.html>.
100. Taieb SB, Taylor JW, Hyndman RJ. Coherent Probabilistic Forecasts for Hierarchical Time Series. In: Precup D, Teh YW, editors. Proceedings of the 34th International Conference on Machine Learning. vol. 70 of Proceedings of Machine Learning Research. PMLR; 2017. p. 3348–3357. Available from: <https://proceedings.mlr.press/v70/taieb17a.html>.