

The Careless Coupon Collector’s Problem

Emilio Cruciani   

European University of Rome, Italy

Aditi Dudeja  

The School of Data Science, The Chinese University of Hong Kong, Shenzhen, China

Abstract

We initiate the study of the Careless Coupon Collector’s Problem (CCCP), a novel variation of the classical coupon collector, that we envision as a model for information systems such as web crawlers, dynamic caches, and fault-resilient networks. In CCCP, a collector attempts to gather n distinct coupon types by obtaining one coupon type uniformly at random in each discrete round, however the collector is *careless*: at the end of each round, each collected coupon type is independently lost with probability p . We analyze the number of rounds required to complete the collection as a function of n and p . In particular, we show that it transitions from $\Theta(n \ln n)$ when $p = o(\frac{\ln n}{n^2})$ up to $\Theta((\frac{np}{1-p})^n)$ when $p = \omega(\frac{1}{n})$ in multiple distinct phases. Interestingly, when $p = \frac{c}{n}$, the process remains in a metastable phase, where the fraction of collected coupon types is concentrated around $\frac{1}{1+c}$ with probability $1 - o(1)$, for a time window of length $e^{\Theta(n)}$. Finally, we give an algorithm that computes the expected completion time of CCCP in $O(n^2)$ time.

2012 ACM Subject Classification Mathematics of computing

Keywords and phrases Coupon Collector, Markov Chains, Metastability

Digital Object Identifier 10.4230/LIPIcs.FUN.2026.14

Related Version *Full Version*: <https://arxiv.org/abs/2602.20705>

Funding *Aditi Dudeja*: This work was initiated while the author was affiliated with the University of Salzburg. This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 947702).

Acknowledgements We thank Silvia Cruciani for the CCCP illustration.

1 Introduction

The coupon collector’s problem is a fundamental paradigm in probability theory and combinatorics, with a history dating back to 1708 as a problem formulated by Abraham de Moivre [16] and still studied today (see Section 1.1 for a discussion of related work). In its classical form, a collector seeks to complete a collection of n distinct coupons by receiving one coupon uniformly at random in each discrete round. The natural question is that of understanding the expected number of rounds required to complete the collection. Beyond its theoretical elegance, the model has become a cornerstone in computer science, widely utilized for the analysis of randomized algorithms and data structures, providing an essential tool for understanding properties of various probabilistic processes [26, 24].

In this paper, we initiate the study of a novel variation that we call the Careless Coupon Collector’s Problem (CCCP). Unlike the traditional collector who retains every acquired item, the *careless collector* may lose progress over time. This model serves as a natural abstraction for modern dynamic information systems, such as *web crawlers* maintaining fresh indices, *distributed caches* with randomized expiration policies, and *fault-resilient networks* where sensors may periodically disconnect. The modeling of these systems through CCCP is discussed in more detail in Section 1.2.



© Emilio Cruciani and Aditi Dudeja;
licensed under Creative Commons License CC-BY 4.0

13th International Conference on Fun with Algorithms (FUN 2026).

Editor: John Iacono; Article No. 14; pp. 14:1–14:21



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

14:2 The Careless Coupon Collector's Problem

Formal definition of CCCP

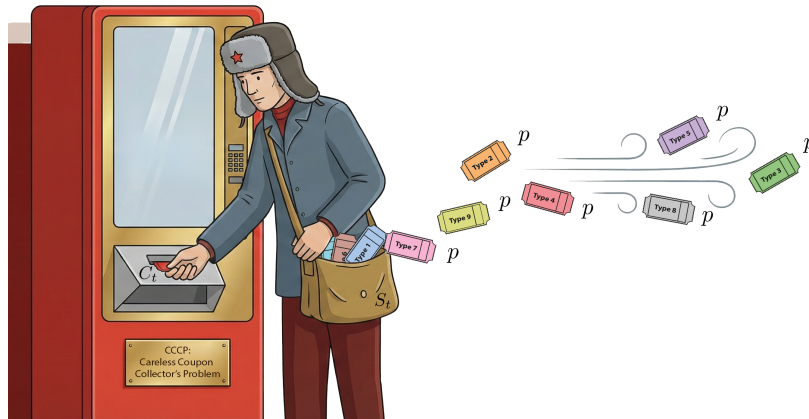
Fix an integer $n \geq 1$, the number of distinct coupon types, and a parameter $p = p(n) \in [0, 1]$, a probability loss parameter. Let $[n] := \{1, 2, \dots, n\}$. We define a discrete-time Markov chain $(S_t)_{t \geq 0}$ with state space $2^{[n]}$ modeling the set of coupons acquired by the careless collector in round t . In particular, starting from $S_0 = \emptyset$, in each round $t \geq 0$ we define

$$S_{t+1} = (S_t \cup \{C_t\}) \setminus \{i : L_{i,t} = 1\},$$

where:

- C_t is a random element of $[n]$ drawn independently and uniformly at random, modeling the new coupon acquired by the collector in round t ;
- $\{L_{i,t}\}_{i \in [n]}$ are independent and identically distributed indicator random variables with parameter p , each modeling if the collector loses coupon i in round t (i.e., $L_{i,t} = 1$) or keeps it (i.e., $L_{i,t} = 0$).

Equivalently, at each step the collector first gains the coupon C_t , then independently loses each coupon in their possession (including C_t) with probability p . Note that obtaining a coupon $C_t \in S_t$ does not have any effect on the collection. Also note that the collector may lose multiple coupons in the same round, including the one they just obtained. See Figure 1 for an illustration of CCCP.



■ **Figure 1** An illustration of a careless coupon collector.

We are interested in the first round t such that the collector acquires all coupon types, namely at the following *hitting time*

$$T_{n,p} := \min\{t \geq 0 : S_t = [n]\}.$$

If $\{t \geq 0 : S_t = [n]\} = \emptyset$ the hitting time is undefined and we say $T_{n,p} = \infty$. Our goal is to study the asymptotic behavior of $\mathbb{E}[T_{n,p}]$ as a function of n and p .


Behavior of the dynamics

As discussed in detail in Section 2.1, the process generalizes the classical coupon collector's problem. However, the loss of coupons makes CCCP significantly harder to analyze, since it breaks its monotonic growth and introduces “catastrophic” downward transitions (see e.g., *skip-free Markov Chains* [15] or *birth-and-death processes with catastrophes* [11]), leading to a more complex behavior.

In particular, the loss parameter p has two main effects on the dynamics. The first is that of slowing down the completion of the collection. As p increases, the expected hitting time undergoes a sharp qualitative change. Intuitively, for small p , the losses are negligible and the process behaves essentially like the classical coupon collector, completing quickly. As p grows, losses start to matter and the hitting time increases rapidly, transitioning through distinct regimes. In particular, small variations of p can trigger large, even exponential, increases in the hitting time, reflecting a phase-like shift from fast collection to a regime dominated by rare events (i.e., lucky streaks of acquisitions that allow for the completion of the collection).

The second effect of p is to induce a metastable phase. Intuitively, the dynamics is governed by two competing forces: the acquisition of new coupons and the loss of already collected ones, both depending on n, p , and on the current number of collected coupons. As the process evolves, these forces balance each other, creating an attractor of the dynamics that traps the system in a set of states where the fraction of collected coupons fluctuates around a limiting value $q_* = q_*(n, p)$.

In Figure 2 we report some simulations illustrating these phenomena. In particular, on the left we see the first effect of p on the hitting time, that quickly transitions from being log-linear to exponential. On the right, we observe how the fraction of collected coupons quickly stabilizes around the limiting fraction q_* . When $p = \Omega(1/n)$, the expected number of coupons nq_* held at any given time becomes very small, dropping to a constant or even sub-constant value. In this regime, the natural state of the system is to be nearly empty. Consequently, completing the collection of all n coupons requires the system to deviate drastically from its equilibrium. Reaching the full set essentially becomes a large stochastic fluctuation. In other words, the collector remains trapped in a metastable state where the number of coupons fluctuates around nq_* for an exponential amount of time before a rare event allows to complete the collection.



■ **Figure 2** Empirical hitting time and metastable fractions of coupons of CCCP, averaged over 1000 repetitions. On the left: average hitting time while varying p , for fixed $n = 10$; the small value of n allows to wait exponential times for the process to end. On the right: average fraction $|S_t|/n$ of coupons in the collection at time t , for several values of p .

Our contributions

Our contributions can be summarized as follows:

- **Model:** We introduce CCCP, a novel and natural generalization of the coupon collector's problem that, apart from the mathematical interest, can be used to model modern dynamic information systems (see Section 1.2).

14:4 The Careless Coupon Collector’s Problem

- **Metastability:** We prove the existence of a *metastable phase*, where the fraction of collected coupons remains concentrated for exponentially-many rounds with probability $1 - o(1)$ around a limiting quantity q_* that we exactly characterize as a function of n, p (see Section 3).
- **Hitting time:** We study the *expected hitting time* at three different levels (see Section 4):
 - We give an efficient *exact algorithm* for the expected hitting time that runs in time $O(n^2)$ (Section 4.1).
 - We perform a *mean-field analysis* of the expected hitting time, through a spatiotemporal independence assumption: we give a simple asymptotic formulation that intuitively explains the hitting time of CCCP as a sum of two terms: the (marginal) *mixing time* to the metastable state, and the *escape time* from the metastable trap, waiting for the rare event in which the collection is completed (Section 4.2). The expected hitting time described by this analysis essentially matches our simulations we presented earlier.
 - We provide rigorous *unconditional upper and lower bounds* to the expected hitting time (Section 4.3). Our lower bound is asymptotically tight, matching the mean-field analysis.

Future research directions

As CCCP is a newly proposed model, the primary significance of this work lies in identifying the collector’s regimes and the discovery of the metastable phase. We believe this provides a foundation for future research directions:

- The most immediate theoretical challenge is to provide a more refined upper bound to the expected hitting time that matches the lower bound, that we believe to be tight as it is asymptotically equal to the mean-field analysis.
- In this work, we assume independent and uniform losses. Future studies could investigate losses where each coupon type has its own probability of being loss p_i , on the same line as other classical models where the sampling of C_t is not uniform, e.g., [17].

1.1 Related work

As already mentioned, the problem’s origins date back to 1708, appearing as a challenge posed by De Moivre [16]. However, the modern era of the coupon collector’s problem began in the mid-20th century with rigorous results regarding hitting times and limit distributions established by von Schelling [29], Newman [25], and the seminal work of Erdős and Rényi [14]. The theoretical framework was later expanded by viewing the process through different lenses: Holst [19] framed the collector’s progress as a classic occupancy problem in urn models; Flajolet et al. [17] utilized the context of formal languages and generating functions to analyze birthday and coupon collection paradoxes.

Due to its elegant simplicity, the coupon collector’s problem has found utility across diverse sectors. In engineering it has been used to model reliability and optimal testing strategies [23, 12]. In biology to model probabilistic structures in biological systems [1]. Moreover, it is a core concept in the analysis of randomized algorithms [26, 24] and has been the subject of extensive surveys regarding optimization [10].

Recent research has shifted toward specialized variations that relax classical assumptions. Under arbitrary sampling distributions, rather than uniform, Berenbrink and Sauerwald [8] investigated approximations of the hitting times while Anceaume et al. [5, 4] studied the time required to collect a given number of distinct coupons. Raz and Zhan [27] explored a memory-bounded variant where the collector must decide when to stop based on limited

space, providing insights into the random-query model of computation. Other studies have expanded the model to include the collection of multiple copies of each coupon [2], collecting coupons “with friends” [3], competing with other players [21], through the best-of-the-sample paradigm where one samples a set and keeps the rarest items [30], and through the notion of super-coupons [7]. The problem continues to be generalized into highly abstract domains, such as quantum information and graph theory. The introduction of the quantum coupon collector explores sampling in a quantum state framework [6]. The labeled coupon collector has been applied to recovering bipartite graphs with perfect matchings [9].

Closer to the spirit of this paper, there exist also variants of the coupon collector that consider adversarial elements like resets, where the collector faces penalties or can lose all its coupons during the collection [28, 20].

1.2 Applications

This section highlights how analyzing key properties of CCCP, such as the expected hitting time and typical metastable states, can inform the design and analysis of dynamic information systems, including web crawlers, distributed caches, and fault-resilient networks.

Dynamic data under web crawling

A web crawler monitors n web pages that continuously change over time. After visiting a page, its cached copy may become outdated if the page is modified. We model this by assuming that each cached page becomes stale independently with probability p at every step, a parameter driving the rate of content change. The crawler revisits one page at a time, replacing stale copies with fresh ones. The hitting time $T_{n,p}$ is the first moment when all n cached pages are simultaneously up to date. Only when p is sufficiently small all pages align as fresh at once, corresponding to a complete, system-wide refresh of the index. Most of the time, the crawler operates in a metastable state, where a stable fraction q_* of pages remain fresh, while the rest lag behind. The CCCP analysis can be used to estimate how frequently the crawler must revisit pages to maintain a desired level of global freshness, since higher revisit frequency corresponds to smaller p and thus higher overall consistency.

Cache expiration under load-balancing policies

A distributed cache maintains n items replicated across servers using a randomized expiration policy for load balancing. To prevent cache hotspots or synchronization bursts, it periodically invalidate items at random, with each cached entry that expires independently with probability p at every step. When an item expires, it must be re-fetched from the origin server, similarly to re-collecting a lost coupon. The hitting time $T_{n,p}$ represents the time until all n items are simultaneously valid in the cache. In practice, this full synchronization occurs rarely when p is moderately large, while the system typically remains in a metastable regime where a stable fraction q_* of cache entries are valid and the rest are pending refresh. The CCCP analysis can help determine how frequently updates should occur to maintain high consistency across replicas, since more frequent updates correspond to smaller p and thus shorter expected refresh times.

Fault resilience in distributed networks

Consider a network of n sensors or devices that periodically report their status to a central server. Each device may stop reporting with probability p at each step, e.g., due to disconnection or failure, and the server polls one device at a time to re-establish its connection.

14:6 The Careless Coupon Collector's Problem

The system is fully synchronized only when all devices have recently reported, a rare event analogous to the hitting time $T_{n,p}$. Most of the time, the system stays in a metastable regime where a steady fraction q_* of devices are active and up to date. The CCCP analysis helps to determine how frequently the monitoring should occur to maintain high system observability and minimize missing reports.

2 Preliminaries

We start the technical part by giving some simple observations regarding CCCP in Section 2.1. Subsequently, in Section 2.2 we analyze the marginal probability of a coupon being present in the collected set of coupons at any time t . The marginal probability analysis will be our key tool in the remainder of the paper.

2.1 Some observations on CCCP

We discuss how $\mathbb{E}[T_{n,p}]$ varies depending on different values of p .

► **Observation 1.** For $p = 0$, CCCP coincides with the classical coupon collector's problem, for which

$$\mathbb{E}[T_{n,0}] = n \cdot H_n,$$

where $H_n = \sum_{i=1}^n \frac{1}{i} = \ln n + O(1)$ denotes the n -th harmonic number (see, for example, [24, Section 2.4.1]). Moreover, by standard probabilistic arguments, for any $c > 0$ it holds that

$$\Pr(T_{n,0} > \mathbb{E}[T_{n,0}] + cn) = o(1).$$

It's obvious that $T_{n,p}$ is stochastically increasing in p , as having a higher chance of losing coupons cannot make the collection any faster. In particular, the following stochastic domination (that we indicate with \preceq) holds true.

► **Proposition 2.** Let $(S_t^{(p_1)})_{t \geq 0}$ and $(S_t^{(p_2)})_{t \geq 0}$ be CCCP chains with loss probabilities p_1 and p_2 . If $p_1 \leq p_2$ then $T_{n,p_1} \preceq T_{n,p_2}$.

Proof. We prove the statement by a monotone coupling argument. We use the same sampled coupon $C_t \in [n]$ for both chains at each time t . For each coupon $i \in [n]$ and time $t \geq 0$, generate a uniform random variable $U_{i,t} \sim \text{Uniform}[0, 1]$, and set

$$L_{i,t}^{(p_1)} = \mathbf{1}\{U_{i,t} \leq p_1\} \quad \text{and} \quad L_{i,t}^{(p_2)} = \mathbf{1}\{U_{i,t} \leq p_2\},$$

where $\mathbf{1}\{E\}$ is the indicator random variable of the event E . Since $p_1 \leq p_2$, then $L_{i,t}^{(p_1)} \leq L_{i,t}^{(p_2)}$ for all i, t . By induction on t , for all $t \geq 0$ we have that $S_t^{(p_1)} \supseteq S_t^{(p_2)}$. Indeed, $S_0 = \emptyset$ in CCCP and adding the same drawn coupon C_t and removing fewer lost coupons for p_1 preserves the inclusion. Since $T_{n,p} := \min\{t \geq 0 : S_t = [n]\}$, it follows immediately that $T_{n,p_1} \preceq T_{n,p_2}$. ◀

We now state some other observations about $T_{n,p}$ for $p > 0$.

► **Observation 3.** For $p = 1$, CCCP never ends since in each time step a coupon is collected and also lost with probability 1. Hence the process never reaches $[n]$ and we have

$$T_{n,1} = \infty.$$

► **Observation 4.** For $p < 1$, CCCP eventually completes the collection with probability 1, hence

$$T_{n,p} < \infty.$$

This holds because, from any state S , there is a positive probability of reaching the target state $[n]$ in a finite number of steps. Indeed there is a probability $(1 - \frac{|S|}{n})(1 - p)^{|S|+1}$ of acquiring a new coupon without losing any existing ones, making a step forward towards $[n]$. Since this probability is strictly positive $T_{n,p}$ is finite.

2.2 Marginal probabilities

We analyze the probability of having a specific coupon in the collection at time t . For each fixed coupon type i , let

$$q_t = q_{i,t} := \Pr(i \in S_t). \quad (1)$$

Note that by symmetry $q_{i,t} = q_t$ for every i , hence we drop the i in the notation.

► **Lemma 5.** Let

$$a := (1 - p)\left(1 - \frac{1}{n}\right) \quad \text{and} \quad b := \frac{1 - p}{n}.$$

For every t , it holds that

$$q_{t+1} = aq_t + b. \quad (2)$$

Proof. By the law of total probability, we have

$$\begin{aligned} q_{t+1} &= \Pr(i \in S_{t+1}) \\ &= \Pr(i \in S_{t+1} \wedge i \in S_t) + \Pr(i \in S_{t+1} \wedge i \notin S_t) \\ &= \Pr(i \in S_{t+1} \mid i \in S_t) \cdot \Pr(i \in S_t) + \Pr(i \in S_{t+1} \mid i \notin S_t) \cdot \Pr(i \notin S_t) \\ &= (1 - p)q_t + \frac{(1 - p)}{n}(1 - q_t). \end{aligned}$$

By definition $a := (1 - p)\left(1 - \frac{1}{n}\right)$ and $b := \frac{1 - p}{n}$. Thus, we conclude that $q_{t+1} = aq_t + b$. ◀

► **Lemma 6.** The recurrence in Equation (2) converges to a unique fixed point

$$q_* = \lim_{t \rightarrow \infty} q_t = \frac{1 - p}{1 - p + np}. \quad (3)$$

Proof. Let $f(x) = ax + b$ be the mapping defining the recurrence in Equation (2). Since $|a| = (1 - p)\left(1 - \frac{1}{n}\right) < 1$, for every x, y we have $|f(x) - f(y)| < |x - y|$, i.e., f is a contraction mapping. By the Banach fixed-point theorem, f admits a unique fixed point q_* . Consequently, $\lim_{t \rightarrow \infty} q_t = q_*$ regardless of the initial value q_0 . Solving the fixed-point equation $q_* = aq_* + b$ yields $q_* = \frac{b}{1 - a} = \frac{1 - p}{1 - p + np}$. ◀

The convergence of q_t to the unique fixed point q_* can be understood as the system reaching a *stochastic steady state*. This state represents the point where the two fundamental forces of CCCP, i.e., coupon acquisition and independent loss, reach an equilibrium. When the size of the collection is less than nq_* the acquisition rate is higher than the cumulative loss rate, and viceversa. The fixed point q_* thus serves as an attractor where the expected net change in the collection size is zero, defining a level around which the process will fluctuate for the majority of its duration, as we will see in the next section, before eventually hitting the completion state $[n]$.

14:8 The Careless Coupon Collector's Problem

► **Corollary 7.** Recall the definition of q_* in Equation (3). It holds that

$$q_* = \begin{cases} 1 - \Theta(np) & \text{if } p = o(\frac{1}{n}), \\ \frac{1}{1+c} - o(1) & \text{if } p = \frac{c}{n}, \text{ for constant } c > 0, \\ \Theta(\frac{1}{np}) & \text{if } p = \omega(\frac{1}{n}). \end{cases}$$

Proof. We show the proof for each of the regimes.

1. Consider the case when $p = o(\frac{1}{n})$. In this case, we use Taylor's expansion $\frac{1}{1+x} = 1 - x + O(x^2)$ where $x = \frac{1}{1+(np-p)}$. Using routine algebra, we can simplify to get the desired bound.
2. When $p = \frac{c}{n}$, we can simplify the expression for q_* to obtain, $q_* = \frac{1-p}{1+c-p}$. This can be further rewritten as $q_* = \frac{1}{1+c} - o(1)$.
3. Consider the case when $p = \omega(\frac{1}{n})$. Then, $q_* = (\frac{1}{np}) \cdot \frac{1-p}{1+\frac{1-p}{np}}$. In this regime, we can rewrite $np = o(1)$. Thus, $q_* = \Theta(\frac{1}{np})$. ◀

► **Lemma 8.** Consider the recurrence in Equation (2) with initial condition q_0 and fixed point q_* as defined in Equation (3). It holds that

$$q_t - q_* = a^t(q_0 - q_*).$$

As a corollary:

- If $q_0 < q_*$, then q_t is monotonically increasing.
- If $q_0 > q_*$, then q_t is monotonically decreasing.

Proof. Since q_* is a fixed point, it holds that $q_* = aq_* + b$. Therefore

$$\begin{aligned} q_t - q_* &= aq_{t-1} + b - (aq_* + b) \\ &= a(q_{t-1} - q_*) \\ &\vdots \\ &= a^t(q_0 - q_*). \end{aligned}$$

Lemma 6 states that q_t converges to q_* and Lemma 8 provides its convergence rate. We also define the *marginal mixing time* as the first time such that q_t is sufficiently close to q_* .

► **Definition 9 (Marginal mixing time).** For any $\epsilon \in (0, 1)$, let

$$\hat{T}_{\text{mix}} = \hat{T}_{\text{mix}}(\epsilon) := \min \{ t \geq 0 : \forall s \geq t, |q_s - q_*| \leq \epsilon \}. \quad (4)$$

Note that the marginal mixing time differs from the standard notion of mixing time [22] since it only measures convergence of the marginals rather than of the entire state distribution.

► **Lemma 10.** For any $\epsilon \in (0, 1)$, it holds that

$$\hat{T}_{\text{mix}}(\epsilon) = \frac{\ln(q_*/\epsilon)}{\ln(1/a)} = \Theta\left(\frac{\ln(q_*/\epsilon)}{p + 1/n}\right).$$

Proof. By Lemma 8, for any arbitrary small $\epsilon > 0$, we have that

$$|q_t - q_*| = a^t q_* \leq \epsilon \iff t \geq \frac{\ln(q_*/\epsilon)}{\ln(1/a)} = \hat{T}_{\text{mix}}.$$

Using the Taylor expansion $\ln(\frac{1}{1-x}) = x(1 + o(1))$, since $a := (1-p)(1-1/n)$ we have $\ln(1/a) = (p + \frac{1}{n})(1 + o(1))$. Therefore, we conclude that $\hat{T}_{\text{mix}}(\epsilon) = \Theta(\frac{\ln(q_*/\epsilon)}{p+1/n})$. ◀

3 Metastability

For every time $t \geq 0$ and coupon $i \in [n]$, let

$$X_{i,t} := \mathbf{1}\{i \in S_t\}, \quad (5)$$

namely $X_{i,t}$ is an indicator random variable of the event $\{i \in S_t\}$ that is 1 with probability q_t , by Equation (1). Hence, it holds that $\mathbb{E}[X_{i,t}] = q_t$.

► **Lemma 11.** *It holds that*

$$\mathbb{E}[|S_t|] = nq_t.$$

Proof. For each $t \geq 0$, by linearity of expectation we get

$$\mathbb{E}[|S_t|] = \mathbb{E}\left[\sum_{i=1}^n X_{i,t}\right] = \sum_{i=1}^n \mathbb{E}[X_{i,t}] = nq_t. \quad \blacktriangleleft$$

For each fixed time t , the random variables $(X_{i,t})_{i \in [n]}$ are not independent, however they are *negatively associated*, as we prove next.

► **Definition 12** (Negative association). *The random variables X_1, \dots, X_n are said to be negatively associated if for all disjoint subsets $I, J \subseteq \{1, \dots, n\}$ and all non-decreasing functions f and g it holds that*

$$\mathbb{E}[f(X_i, i \in I) \cdot g(X_j, j \in J)] \leq \mathbb{E}[f(X_i, i \in I)] \cdot \mathbb{E}[g(X_j, j \in J)].$$

► **Lemma 13.** *For every fixed t , it holds that $(X_{i,t})_{i \in [n]}$ are negatively associated.*

Proof. Recall that $S_0 = \emptyset$ and S_t depends only on S_{t-1} , on the sampled coupon C_{t-1} and indicators $\{L_{i,t-1}\}_{i \in [n]}$. In particular, whether coupon i belongs to S_t depends only on coupons drawn at times $s \in \{0, \dots, t-1\}$ and whether those tokens survived the sequence of losses after up to time $t-1$. Formally, for each time $s \in \{0, \dots, t-1\}$ and coupon $i \in [n]$ define

$$Y_{i,s} := Z_{i,s} \cdot R_{i,s},$$

where

$$Z_{i,s} := \mathbf{1}\{C_s = i\} \quad \text{and} \quad R_{i,s} := \prod_{r=s}^{t-1} \mathbf{1}\{L_{i,r} = 0\}$$

respectively indicate that coupon i was sampled at time s , and that specific token survived every subsequent loss event up to time $t-1$. With this definition, it holds that

$$X_{i,t} = \mathbf{1}\{i \in S_t\} = \bigvee_{s=0}^{t-1} Y_{i,s},$$

namely coupon i is in S_t if and only if there exists at least one draw before time t that sampled coupon i and coupon i survived all later losses up to $t-1$.

We prove negative association by starting from a family of negatively associated random variables and leveraging closure properties of negatively associated random variables [13, Chapter 3.1]. For each fixed s , the family of indicators $Z_{i,s}$ are such that exactly one variable

14:10 The Careless Coupon Collector's Problem

is 1 and the rest are 0 and therefore is negatively associated. For each fixed s , the $R_{i,s}$ are independent across i and independent of the $Z_{i,s}$. Combining a negatively associated family with an independent family preserves negative association. Moreover the families $(Z_{i,s}, R_{i,s})$ are independent across different s . Similarly, concatenating two independent families, each negatively associated, preserves negative association. Finally, each $X_{i,t}$ is a coordinate-wise non-decreasing function of a disjoint block of $Y_{i,s}$'s. Therefore, by the monotone aggregation property, the family $(X_{i,t})_i$ is negatively associated. \blacktriangleleft

Given the negative dependence, we can apply the following version of the Chernoff bound to prove that the expected fraction of collected coupons stays concentrated around q_* for an exponential number of rounds.

► **Theorem 14** (Multiplicative Chernoff bound, [13]). *Let X_1, \dots, X_n be independent or negatively associated indicator random variables. Let $X := \sum_{i=1}^n X_i$ and let $\mu = \mathbb{E}[X]$. For all $\delta \in (0, 1)$ it holds that:*

$$\Pr(|X - \mu| > \delta\mu) \leq 2e^{-\frac{\delta^2}{3}\mu}.$$

In particular, with the following proposition we prove the metastable regime for small values, namely when $p = O(1/n)$.

► **Proposition 15** (Metastable number of coupons). *For any $\delta \in (0, 1)$ pick a positive $\epsilon \leq \delta q_*$ and let $T = \hat{T}_{\text{mix}}(\epsilon)$. For every $L \in \mathbb{N}$, it holds that*

$$\Pr(\exists t \in \{T, \dots, T + L - 1\} : |S_t| - nq_* > 2\delta nq_*) \leq 2Le^{-\frac{\delta^2(1-\delta)}{3}nq_*}.$$

As a corollary, if $p = O(1/n)$, there exists $L = e^{\Theta(n)}$ such that

$$\Pr(\exists t \in \{T, \dots, T + L - 1\} : |S_t| - nq_* > 2\delta nq_*) = o(1).$$

Proof. Note that, by triangle inequality, we have

$$\begin{aligned} |S_t| - nq_* &= |S_t - nq_t + nq_t - nq_*| \\ &\leq |S_t - nq_t| + |nq_t - nq_*| \\ &= |S_t - nq_t| + n|q_t - q_*|. \end{aligned}$$

We analyze the two terms in the previous bound separately. For the first term, note that, by Definition 9, for every $t \geq T = \hat{T}_{\text{mix}}(\epsilon)$ we have that $|q_t - q_*| \leq \epsilon$, which implies $q_t \geq q_* - \epsilon \geq (1 - \delta)q_*$, since $\epsilon \leq \delta q_*$ by assumption. Hence, since $\mathbb{E}[|S_t|] = nq_t$ by Lemma 11, for every fixed $t \geq T$ by Theorem 14 we have

$$\Pr(|S_t| - nq_t > \delta nq_t) \leq 2e^{-\frac{\delta^2}{3}nq_t} \leq 2e^{-\frac{\delta^2(1-\delta)}{3}nq_*}.$$

For the second term, by Definition 9, for every $t \geq T$ deterministically we have

$$n|q_t - q_*| \leq n\epsilon \leq \delta nq_*.$$

Therefore, for every fixed $t \geq T$ we have

$$\Pr(|S_t| - nq_* > 2\delta nq_*) \leq 2e^{-\frac{\delta^2(1-\delta)}{3}nq_*}.$$

By taking a union bound over all times $t \in \{T, \dots, T + L - 1\}$ we conclude that

$$\Pr(\exists t \in \{T, \dots, T + L - 1\} : |S_t| - nq_* > 2\delta nq_*) \leq 2Le^{-\frac{\delta^2(1-\delta)}{3}nq_*}.$$

Regarding the corollary, by Corollary 7, whenever $p = O(1/n)$ we have $q_* = \Omega(1)$. Hence the corollary follows by picking, $L = e^{cn}$ for any $c < \delta^2(1 - \delta)q_*/3$. \blacktriangleleft

When the mean is small, or p is large, we cannot get as strong guarantees as implied by the previous proposition. Instead, we show the following bound.

► **Proposition 16.** *Assume $p = \omega(1/n)$ and let $T = \hat{T}_{\text{mix}}(\epsilon)$. For any $\delta \in (0, 1)$ and every $L \in \mathbb{N}$, it holds that*

$$\Pr(\exists t \in \{T, \dots, T + L - 1\} : |S_t| > \delta n) \leq Le^{-\frac{\delta^2}{4}n}.$$

As a corollary, there exists $L = e^{\Theta(n)}$ such that

$$\Pr(\exists t \in \{T, \dots, T + L - 1\} : |S_t| > \delta n) = o(1).$$

The previous proposition covers the remaining case $p = \omega(1/n)$. We note that the absolute deviations from the mean in Propositions 15 and 16 have comparable magnitude, as $q_* = \Theta(1)$ for $p = O(1/n)$. To prove Proposition 16, we use the following version of Chernoff.

► **Theorem 17** (Additive Chernoff bound, [13]). *Let X_1, \dots, X_n be independent or negatively associated indicator random variables. Let $X := \sum_{i=1}^n X_i$ and let $\mu = \mathbb{E}[X]$. For any $\Delta > 0$ it holds that:*

$$\Pr(X - \mu > \Delta) \leq e^{-\frac{2\Delta^2}{n}}.$$

Proof of Proposition 16. By Lemma 11, the fact that $q_t \leq q_*$ for all t , and Corollary 7, since $p = \omega(1/n)$ it follows that $\mathbb{E}[|S_t|] = nq_t \leq nq_* = o(n)$. Therefore, for any constant $\delta \in (0, 1)$, it holds that $\delta n - q_t n \geq \frac{\delta}{2}n$. By Theorem 17 we have

$$\Pr(|S_t| > \delta n) = \Pr(|S_t| - q_t n > \delta n - q_t n) \leq e^{-\frac{(\delta n - q_t n)^2}{n}} = e^{-(\delta - q_t)^2 n} \leq e^{-\frac{\delta^2}{4}n}.$$

By taking a union bound over all times $t \in \{T, \dots, T + L - 1\}$ we conclude that

$$\Pr(\exists t \in \{T, \dots, T + L - 1\} : |S_t| > \delta n) \leq Le^{-\frac{\delta^2}{4}n}.$$

The corollary follows by picking $L = e^{cn}$, for any $c < \delta^2/4$. ◀

4 Hitting time

In this section we study the expected hitting time of CCCP. In particular in Section 4.1 we give an algorithm to compute the exact hitting time, in Section 4.2 we perform a mean-field analysis of the expected hitting time, and in Section 4.3 we give upper and lower bounds that drop the mean-field assumptions.

4.1 Exact computation

In this section we start to study the expected hitting time of CCCP. As we will see, the exact formulation does not admit a closed-form expression. However, we describe how it can be computed exactly, proving the following statement.

► **Proposition 18.** *There is an algorithm that, given input n, p , computes $\mathbb{E}[T_{n,p}]$ in time $O(n^2)$.*

Proof. Let $K_t = |S_t|$. Note that in order to compute the expected hitting time $\mathbb{E}[T_{n,p}]$ we can study the simpler Markov Chain $(K_t)_{t \geq 0}$ associated to CCCP and note that


$$T_{n,p} = \min\{t \geq 0 : K_t = n\}.$$

14:12 The Careless Coupon Collector's Problem

is an equivalent formulation of the hitting time of CCCP. This process is simpler in that it consists only of $n + 1$ states, i.e., $\{0, 1, \dots, n\}$. In particular, the transition matrix $P \in \mathbb{R}^{(n+1) \times (n+1)}$ of the reduced process can be fully characterized as follows. For every k , we have:

$$\begin{cases} P_{k,k-i} = \left(1 - \frac{k}{n}\right) \binom{k+1}{i+1} (1-p)^{k-i} p^{i+1} + \frac{k}{n} \binom{k}{i} (1-p)^{k-i} p^i, & \forall i \in \{0, \dots, k\}, \\ P_{k,k+1} = \left(1 - \frac{k}{n}\right) (1-p)^{k+1}, \\ P_{k,k+i} = 0, \end{cases} \quad \forall i \in \{2, \dots, n-k\}.$$

Note that P is stochastic (rows sums to 1) and has a special structure, being a *lower-Hessenberg matrix*, i.e., all entries above the first superdiagonal are 0 [18]. We give a graphical representation of CCCP and its corresponding transition matrix in Figure 3. Note that we will compute $\mathbb{E}[T_{n,p}]$ by solving a system of linear equations described by the transition matrix of CCCP. We mention the special structure of this matrix because it will be leveraged to speed-up the computation of $\mathbb{E}[T_{n,p}]$.



■ **Figure 3** A graphical representation of the reduced Markov Chain of CCCP with $n = 3$ coupons (left) and its corresponding lower-Hessenberg transition matrix with first superdiagonal entries in magenta (right).

In order to do so, let us define now

$$h(k) := \mathbb{E}[T_{n,p} \mid K_0 = k].$$

Since in CCCP $K_0 = 0$, by the law of total expectation we compute $h(0) = \mathbb{E}[T_{n,p} \mid K_0 = 0] = \mathbb{E}[T_{n,p}]$ recursively as follows:

$$\begin{cases} h(n) = 0, \\ h(k) = 1 + \sum_{i=0}^{k+1} P_{k,i} \cdot h(i), \quad \forall k \in \{0, \dots, n-1\}, \end{cases}$$

where $h(n) = 0$ is the boundary condition that stops the recursion at the first time the process hits state n , and the sum in $h(k)$ goes up to $k + 1$ since $P_{k,k+i} = 0$ for all $i > 1$. In words, in this recursive formulation we pay 1 for each step of the chain plus the expected hitting time from the state we reach. By plugging in the equivalence $h(n) = 0$ we get a linear system of n equations in n unknowns $h(0), \dots, h(n-1)$. Now, let I be the identity matrix and $Q \in \mathbb{R}^{n \times n}$ be the matrix of transient states of CCCP (i.e., P without the last row and column). We write the system in matrix form and, by rearranging terms, we get

$$A\mathbf{h} = \mathbf{b}, \tag{6}$$

where $A = (I - Q) \in \mathbb{R}^{n \times n}$ is such that

$$A_{k,i} = \begin{cases} 1 - P_{k,i} & \text{if } k = i, \\ -P_{k,i} & \text{if } k \neq i, \end{cases}$$

$\mathbf{b} = (1, \dots, 1)^\top \in \mathbb{R}^n$ is the vector of all ones, and $\mathbf{h} = (h(0), \dots, h(n-1))^\top$ is the vector of unknowns. Note that the matrix A is non-singular since Q is the submatrix of transient states of CCCP and, therefore, the system in Equation (6) has a unique solution that equals the vector of expected absorbing times. Moreover, note that A inherits the lower-Hessenberg structure of P . This structure allows to specialize the Gaussian elimination algorithm, solving the system more efficiently [18]. We proceed as follows:

1. **Forward-elimination:** The goal is to transform the system in Equation (6) into an equivalent one with an upper-triangular coefficient matrix, i.e., all elements below the diagonal are 0. As in standard Gaussian elimination, for each pivot (a diagonal entry of A) we *eliminate* the entries below the pivot and set them to 0 by subtracting the correctly rescaled pivot row, so that the system stays equivalent. As A is lower-Hessenberg, the first row has only two nonzero entries: when we subtract it from all rows below, only $O(n)$ terms are affected instead of $O(n^2)$. After eliminating the first column, also the second row has only two nonzero entries, and so on iteratively for every row. This property allows to perform the forward-elimination step efficiently. More formally, for each $c = 0, \dots, n-2$, for each $r > c$, let $m_{r,c} := A_{r,c}/A_{c,c}$ and perform the following updates:

$$\begin{aligned} A_{r,c} &\leftarrow A_{r,c} - m_{r,c} \cdot A_{c,c} = 0, \\ A_{r,c+1} &\leftarrow A_{r,c+1} - m_{r,c} \cdot A_{c,c+1}, \\ \mathbf{b}_r &\leftarrow \mathbf{b}_r - m_{r,c} \cdot \mathbf{b}_c. \end{aligned}$$

For each c there are $n - c - 1$ entries below the diagonal that need to be eliminated, and each elimination step only updates two entries of A and one entry of \mathbf{b} . Hence, summing over all columns, we have a total running time of $\sum_{c=0}^{n-2} 3(n - c - 1) = O(n^2)$.

2. **Backward-substitution:** After the forward-elimination, the system in Equation (6) has been transformed into an equivalent upper-triangular system

$$U\mathbf{h} = \mathbf{b}'.$$

Note that U is also still lower-Hessenberg, hence it has only two nonzero entries in each row, the diagonal and the first superdiagonal. This also allows to speed-up the backward-substitution step computing the vector \mathbf{h} as follows:

$$\begin{aligned} h(n-1) &= \frac{\mathbf{b}'_{n-1}}{U_{n-1,n-1}}, \\ &\vdots \\ h(k) &= \frac{\mathbf{b}'_k - U_{k,k+1} \cdot h(k+1)}{U_{k,k}}, \\ &\vdots \\ h(0) &= \frac{\mathbf{b}'_0 - U_{0,1} \cdot h(1)}{U_{0,0}}. \end{aligned}$$

Each step costs $O(1)$ to rearrange the equation, so backward-substitution costs $O(n)$ overall.

The procedure computes the entire vector of absorbing states \mathbf{h} and, in particular, the expected hitting time of CCCP $h(0) = \mathbb{E}[T_{n,p}]$ in time $O(n^2) + O(n) = O(n^2)$. ◀

4.2 Mean-field analysis

In this section we perform a mean-field analysis of the expected hitting time $T_{n,p}$. In particular we analyze it by assuming the spatial and temporal independence, as follows. We remark that such assumptions do not hold in CCCP, but for the sake of the argument we make them nonetheless, as they are sufficient to derive clean bounds on the expected hitting time that capture the essential physics of CCCP. Specifically, the analysis reveals a two-stage process: first, a rapid climb where the collection grows quickly toward the metastable state nq_* ; second, a prolonged stochastic drift where the system fluctuates around nq_* for an incredibly long duration. In fact, interestingly, the bounds derived in this section essentially match the empirical values of the hitting time observed in the simulations described in the introduction (see Figure 2).

► **Assumption 19** (Spatial and temporal independence). *Recall the random variables $(X_{i,t})_{i \in [n], t \geq 0}$ defined in Equation (5). Assume:*

■ **Spatial independence:** *For every fixed time t , the random variables $(X_{i,t})_{i \in [n]}$ are independent. Therefore, by Equation (1) we have*

$$\Pr(S_t = [n]) = \Pr\left(\bigwedge_{i=1}^n \{X_{i,t} = 1\}\right) = q_t^n.$$

■ **Temporal independence:** *For every fixed coupon i , the process $(X_{i,t})_{t \geq 0}$ consists of independent random variables across time, and families corresponding to different coupons are mutually independent. Therefore, using also the spatial independence we have*

$$\Pr(T_{n,p} > t) = \Pr\left(\bigwedge_{j=0}^t \{S_j \neq [n]\}\right) = \prod_{j=0}^t (1 - q_j^n).$$

► **Proposition 20** (Mean-field expected hitting time). *Let q_* as in Equation (3) and, for any $\epsilon > 0$, let $\hat{T}_{\text{mix}} = \hat{T}_{\text{mix}}(\epsilon)$ as in Definition 9. Under Assumption 19, it holds that:*

$$\hat{T}_{\text{mix}} \cdot (1 - q_*^n \hat{T}_{\text{mix}}) + \frac{(1 - q_*^n)^{\hat{T}_{\text{mix}} + 1}}{q_*^n} \leq \mathbb{E}[T_{n,p}] \leq \hat{T}_{\text{mix}} + \frac{1}{(q_* - \epsilon)^n}.$$

Proof. For the sake of readability, let $T = \hat{T}_{\text{mix}}(\epsilon)$. Since $T_{n,p} \geq 0$, we can write the expected hitting time using the tail-sum formula as follows:

$$\begin{aligned} \mathbb{E}[T_{n,p}] &= \sum_{t=0}^{\infty} \Pr(T_{n,p} > t) \\ &= \sum_{t=0}^{T-1} \Pr(T_{n,p} > t) + \sum_{t=T}^{\infty} \Pr(T_{n,p} > t) \\ &= \sum_{t=0}^{T-1} \Pr(T_{n,p} > t) + \sum_{k=0}^{\infty} \Pr(T_{n,p} > T + k). \end{aligned} \tag{7}$$

By our choice of T , we can consider the contributions of two terms in Equation (7): first, the *marginal mixing time* where the q_t 's are still small and growing towards q_* (see Lemma 8); second, the *escape time* where the q_t 's are very close to q_* and the hitting time can be approximated by a geometric random variable thanks to the independence assumptions.

In the first term of Equation (7), for $t < T$, we have that $q_t < q_*$. By using that $\Pr(T_{n,p} > t) \leq 1$ we have

$$\sum_{t=0}^{T-1} \Pr(T_{n,p} > t) \leq T.$$

By a union bound and due to the spatial independence assumption, we get

$$\begin{aligned} \sum_{t=0}^{T-1} \Pr(T_{n,p} > t) &= \sum_{t=0}^{T-1} \left(1 - \Pr \left(\bigvee_{j=0}^t \{S_j = [n]\} \right) \right) \\ &\geq \sum_{t=0}^{T-1} \left(1 - \sum_{j=0}^t q_j^n \right) \\ &\geq T(1 - Tq_*^n). \end{aligned}$$

In the second term of Equation (7), for $t \geq T$, we have that $|q_t - q_*| \leq \epsilon$. Therefore, as a consequence of the spatial and temporal independence assumptions we get

$$\begin{aligned} \sum_{k=0}^{\infty} \Pr(T_{n,p} > T+k) &= \sum_{k=0}^{\infty} \prod_{j=0}^{T+k} (1 - q_j^n) \\ &\leq \sum_{k=0}^{\infty} \prod_{j=T}^{T+k} (1 - q_j^n) \\ &\leq \sum_{k=0}^{\infty} (1 - (q_* - \epsilon)^n)^k \\ &= \frac{1}{(q_* - \epsilon)^n}. \end{aligned}$$

Similarly, since $q_0 = 0$ by Lemma 8 we have $q_j \leq q_*$ for all j , and therefore

$$\begin{aligned} \sum_{k=0}^{\infty} \Pr(T_{n,p} > T+k) &= \sum_{k=0}^{\infty} \prod_{j=0}^{T+k} (1 - q_j^n) \\ &\geq \sum_{k=0}^{\infty} (1 - q_*^n)^{T+k+1} \\ &= (1 - q_*^n)^{T+1} \sum_{k=0}^{\infty} (1 - q_*^n)^k \\ &= \frac{(1 - q_*^n)^{T+1}}{q_*^n}. \end{aligned}$$

We can substitute the bounds we derived in Equation (7) to conclude the proof. ◀

► **Observation 21.** *Under Assumption 19, for sufficiently small positive ϵ , we have*

$$\mathbb{E}[T_{n,p}] \approx \hat{T}_{\text{mix}} + q_*^{-n}.$$

14:16 The Careless Coupon Collector's Problem

The previous approximation defines the following *regimes* of the expected hitting time:

$$\mathbb{E}[T_{n,p}] \approx \begin{cases} \Theta(n \ln n) & \text{if } p = o\left(\frac{\ln n}{n^2}\right), \\ n^{\Theta(c)} & \text{if } p = c \frac{\ln n}{n^2}, \text{ for constant } c > 0, \\ e^{\Theta(n^2 p)} & \text{if } p = \omega\left(\frac{\ln n}{n^2}\right) \text{ and } p = o(1/n), \\ \Theta((1+c)^n) & \text{if } p = \frac{c}{n}, \text{ for constant } c > 0, \\ \Theta\left(\left(\frac{np}{1-p}\right)^n\right) & \text{if } p = \omega\left(\frac{1}{n}\right). \end{cases}$$

We describe these regimes in more detail. Towards this, recall the bounds on q_* (Lemma 6) and \hat{T}_{mix} (Lemma 10).

1. **Classical regime:** When $p = o\left(\frac{\ln n}{n^2}\right)$, then \hat{T}_{mix} dominates. One can show that $\hat{T}_{\text{mix}} = \Theta(n \ln n)$. Thus, CCCP asymptotically behaves as if $p = 0$, i.e., as the classical coupon collector's problem.
2. **Super-classical regime:** When $p = \frac{c \ln n}{n^2}$, then q_*^{-n} dominates. One can show that $\mathbb{E}[T_{n,p}]$ is “super-classical” but polynomial. This is because,

$$q_*^{-n} = \left(\frac{1 - \frac{c \ln n}{n^2}}{1 - \frac{c \ln n}{n^2} + \frac{c \ln n}{n}} \right)^{-n} = \left(1 - \frac{\frac{c \ln n}{n}}{1 - \frac{c \ln n}{n^2} + \frac{c \ln n}{n}} \right)^{-n} \approx e^{\frac{c}{1-o(1)} \ln n} = n^{\Theta(c)}.$$

3. **Metastable regime I:** When $p = \omega\left(\frac{\ln n}{n^2}\right)$ and $p = o\left(\frac{1}{n}\right)$, then q_*^{-n} dominates, which we can bound as follows.

$$q_*^{-n} = \left(\frac{1-p}{1-p+np} \right)^{-n} \approx e^{\frac{n^2 p}{1-p}} = e^{\Theta(n^2 p)}.$$

Thus, $\mathbb{E}[T_{n,p}]$ is super-polynomial but sub-exponential.

4. **Metastable regime II:** When $p = \frac{c}{n}$, we can again see that q_*^{-n} dominates. Using routine algebra, we have,

$$q_*^{-n} = \left(1 + \frac{c}{1-p} \right)^n = \Theta((1+c)^n).$$

Thus, $\mathbb{E}[T_{n,p}]$ is exponential.

5. **Metastable regime III:** When $p = \omega\left(\frac{1}{n}\right)$, we can again see that q_*^{-n} dominates. Using routine algebra, we have,

$$q_*^{-n} = \left(\frac{1-p+np}{1-p} \right)^n = \Theta\left(\left(\frac{np}{1-p}\right)^n\right).$$

Thus, $\mathbb{E}[T_{n,p}]$ is super-exponential.

While even more regimes can be defined, we believe these are the most interesting ones: 1 and 2 describe the regime for which the hitting time of CCCP is still polynomial; on the other hand 3, 4, and 5 describe the hitting times in the phases where CCCP has the three distinct regimes of the metastable fraction of collected coupons q_* (see Corollary 7).

Note that in regimes 1, 2, and 3 we have that $q_* = 1 - o(1)$, therefore the metastable window of exponential length described in Propositions 15 and 16 might include fluctuations hitting state $[n]$, not contradicting the sub-exponential hitting time stated above.

Finally, we observe that the concentration of the hitting time $T_{n,p}$ around its expectation varies significantly across these regimes. Recall that in the classical coupon collector's problem ($p = 0$) the hitting time $T_{n,0}$ is concentrated around its expectation (see Observation 1). In the classical regime ($p = o\left(\frac{\ln n}{n^2}\right)$), using Assumption 19, it is possible to prove a similar

concentration of the hitting time $T_{n,p}$. Informally, this is because the probability that the collector completes the collection after the marginal mixing time \hat{T}_{mix} is large enough as the fraction of collected coupon is $q_* = 1 - \Theta(np)$. In contrast, in the other regimes, such a concentration does not hold. Indeed, the expected hitting time is dominated by the escape time from the metastable state which, under our mean-field assumptions, is approximated by a geometric random variable. Since the standard deviation of a geometric distribution is of the same order as its mean, the hitting time $T_{n,p}$ in these regimes undergoes large fluctuations and is not concentrated around its expected value.

4.3 Asymptotic bounds

Finally, in this section we drop Assumption 19 and provide rigorous unconditional upper and lower bounds on the expected hitting time of CCCP. In particular, we first give an upper bound that is still derived by the definition of a geometric random variable but with a different parameter compared to that of the mean-field analysis of Section 4.2. We conclude the section with a lower bound that, instead, is asymptotically tight with respect to our mean-field bound.

Let us consider the reduced chain $(K_t)_{t \geq 0}$ with $K_t = |S_t|$. For $\epsilon \in (0, 1)$, we define

$$\mathcal{G} = \mathcal{G}(\epsilon) := \{k \in \{0, \dots, n\} : k \geq (1 - \epsilon)nq_*\} \quad (8)$$

as the good set of states, namely that set of states that we reach in time \hat{T}_{mix} and from which CCCP hits $[n]$ by waiting for the rare event in which all missing coupons are collected, and the others are not lost.

► **Lemma 22.** *For $k \in \{0, 1, \dots, n\}$ and \mathcal{G} as in Equation (8), it holds that*

$$\min_{k \in \mathcal{G}} \Pr(\exists s \in \{t + 1, \dots, t + n \ln n\} : K_t = k, K_s = n) \geq q_*(1 - \epsilon)(1 - p)^{n^2 \ln n} =: \rho.$$

Proof. For every $m \geq 1$, let

$$F(k, m) := \{\exists s \in \{t + 1, \dots, t + m\} : K_t = k, K_s = n\}.$$

Let us define the following two events needed to satisfy $F(k, m)$:

- E_1 : “all the remaining $n - k$ coupons appear at least once in the m draws of the block”;
- E_2 : “all coupons survive every step in the block”.

Note that $\Pr(F(k, m)) \geq \Pr(E_1 \wedge E_2) = \Pr(E_1) \cdot \Pr(E_2)$ as E_1, E_2 are independent in CCCP (the draws C_t are independent of the losses $L_{i,t}$). We have that

$$\begin{aligned} \Pr(E_1) &\geq 1 - (n - k)(1 - 1/n)^m \geq 1 - (n - k)e^{-m/n}, \\ \Pr(E_2) &\geq (1 - p)^{mn}. \end{aligned}$$

Note that the minimum is achieved for $k = \min(\mathcal{G}) \geq (1 - \epsilon)nq_*$, hence

$$\begin{aligned} \min_{k \in \mathcal{G}} \Pr(F(k, m)) &\geq (1 - (n - k)e^{-m/n}) \cdot (1 - p)^{mn} \\ &\geq (1 - (1 - q_*(1 - \epsilon))ne^{-m/n}) \cdot (1 - p)^{mn}. \end{aligned}$$

The lower bound is (approximately) optimized by picking $m = n \ln n$. Hence, we get

$$\min_{k \in \mathcal{G}} \Pr(F(k, m)) \geq q_*(1 - \epsilon)(1 - p)^{n^2 \ln n}. \quad \blacktriangleleft$$

14:18 The Careless Coupon Collector's Problem

► **Proposition 23.** *Recall the definition of ρ in Lemma 22. It holds that*

$$\mathbb{E}[T_{n,p}] \leq \hat{T}_{\text{mix}}(\epsilon) + \frac{n \ln n}{\rho}.$$

Proof. Recall the definition of \mathcal{G} in Equation (8) and define the bad event E of exiting \mathcal{G} before completing the collection as

$$E := \{ \exists s, \hat{T}_{\text{mix}}(\epsilon) \leq s < T_{n,p} : K_s \notin \mathcal{G} \}.$$

Let \bar{E} be the complement of E . We write the expectation of the hitting time as follows:

$$\begin{aligned} \mathbb{E}[T_{n,p}] &= \mathbb{E}[T_{n,p} \mid E] \cdot \Pr(E) + \mathbb{E}[T_{n,p} \mid \bar{E}] \cdot (1 - \Pr(E)) \\ &\leq \Pr(E) \cdot \mathbb{E}[T_{n,p}] + (1 - \Pr(E)) \cdot \mathbb{E}[T_{n,p} \mid \bar{E}], \end{aligned}$$

where $\mathbb{E}[T_{n,p} \mid E] \leq \mathbb{E}[T_{n,p}]$ holds since the expected hitting time starting from any state not in \mathcal{G} is largest when starting from state 0. By rearranging the terms, we have

$$\mathbb{E}[T_{n,p}] - \Pr(E) \cdot \mathbb{E}[T_{n,p}] \leq (1 - \Pr(E)) \cdot \mathbb{E}[T_{n,p} \mid \bar{E}]$$

and since $\Pr(E) < 1$ we conclude

$$\mathbb{E}[T_{n,p}] \leq \mathbb{E}[T_{n,p} \mid \bar{E}].$$

Let $m = n \ln n$. In the remainder of the proof we show that $\mathbb{E}[T_{n,p} \mid \bar{E}] \leq \hat{T}_{\text{mix}}(\epsilon) + \frac{m}{\rho}$, for ρ as in Lemma 22. We partition the time after $T = \hat{T}_{\text{mix}}(\epsilon)$ into disjoint blocks of length m , where each block is

$$B_i := \{T + im, \dots, T + (i + 1)m - 1\}, \forall i \geq 0.$$

For any $r \geq 0$, define the event

$$F_r := \{T_{n,p} > T + rm - 1\},$$

namely “no block B_0, \dots, B_{r-1} contains a time t where $K_t = n$ ”. Note that

$$\Pr(F_r \mid \bar{E}) \leq (1 - \rho)^r.$$

Recall Equation (7) and write:

$$\begin{aligned} \mathbb{E}[T_{n,p} \mid \bar{E}] &= \sum_{t=0}^{T-1} \Pr(T_{n,p} > t \mid \bar{E}) + \sum_{t=T}^{\infty} \Pr(T_{n,p} > t \mid \bar{E}) \\ &= \sum_{t=0}^{T-1} \Pr(T_{n,p} > t \mid \bar{E}) + \sum_{r=0}^{\infty} \sum_{t \in B_r} \Pr(T_{n,p} > t \mid \bar{E}). \end{aligned}$$

By using $\Pr(T_{n,p} > t \mid \bar{E}) \leq 1$ and noting that $\Pr(T_{n,p} > t)$ is decreasing in t we get

$$\begin{aligned} \mathbb{E}[T_{n,p} \mid \bar{E}] &\leq T + \sum_{r=0}^{\infty} m \Pr(T_{n,p} > T + rm - 1 \mid \bar{E}) \\ &= T + m \sum_{r=0}^{\infty} \Pr(F_r \mid \bar{E}) \\ &\leq T + m \sum_{r=0}^{\infty} (1 - \rho)^r \\ &= T + \frac{m}{\rho}. \end{aligned}$$

We conclude the proof having $\mathbb{E}[T_{n,p}] \leq \mathbb{E}[T_{n,p} \mid \bar{E}] \leq \hat{T}_{\text{mix}}(\epsilon) + \frac{n \ln n}{\rho}$, for ρ as in Lemma 22. ◀

► **Proposition 24.** *It holds that $\mathbb{E}[T_{n,p}] \geq \frac{1}{5q_*^n}$.*

Proof. Recall the definition of $X_{i,t}$ in Equation (5) and that they are negatively associated (Lemma 13). By using negative association and that $q_t \leq q_*$ for all t (by Lemma 8, since $q_0 = 0$), we have

$$\Pr(S_t = [n]) = \Pr\left(\bigwedge_{i=1}^n \{i \in S_t\}\right) \leq \prod_{i=1}^n \Pr(i \in S_t) \leq q_*^n.$$

Moreover, by a union bound

$$\Pr(T_{n,p} \leq t) = \Pr(\exists j \leq t : S_j = [n]) \leq \sum_{j=0}^t \Pr(S_j = [n]) \leq (t+1)q_*^n$$

and therefore

$$\Pr(T_{n,p} > t) \geq 1 - (t+1)q_*^n.$$

Let $\bar{t} := \lfloor \frac{1-2q_*^n}{2q_*^n} \rfloor$ and note that for every $t \leq \bar{t}$ we have

$$\Pr(T_{n,p} > t) \geq \frac{1}{2}.$$

Using the tail-sum formula of the expectation and the previous observation we get

$$\mathbb{E}[T_{n,p}] = \sum_{t=0}^{\infty} \Pr(T_{n,p} > t) \geq \sum_{t=0}^{\bar{t}} \Pr(T_{n,p} > t) \geq \frac{\bar{t}}{2} \geq \frac{1}{5q_*^n}. \quad \blacktriangleleft$$

We acknowledge that while our results correctly identify the exponential nature of the hitting time, there is a significant quantitative gap between the lower and upper bounds in the non-classical regimes. Our lower bound of $\approx 1/q_*^n$ is expected to be tight, matching the mean-field analysis. However, our rigorous upper bound is $\approx 1/(q_*(1-p)^{n^2 \ln n})$. This gap stems from the conservative nature of our block-partitioning argument: to ensure a rigorous result without assuming temporal independence, we bound the probability of completing the set within a safe window where no losses occur. We conjecture that the true hitting time is much closer to the lower bound. Moving from this block-based approach to a more sophisticated analysis of the Markov chain is key for improving the bound.

References

- 1 Ilan Adler, Hyun-Soo Ahn, Richard M Karp, and Sheldon M Ross. A probabilistic model for the survivability of cells. *Journal of Applied Probability*, 42(4):919–931, 2005. URL: <https://www.jstor.org/stable/30040818>.
- 2 Ilan Adler, Shmuel Oren, and Sheldon M Ross. The coupon-collector’s problem revisited. *Journal of Applied Probability*, 40(2):513–518, 2003. URL: <https://www.jstor.org/stable/3215807>.
- 3 Dan Alistarh and Peter Davies. Collecting coupons is faster with friends. In *International Colloquium on Structural Information and Communication Complexity*, pages 3–12. Springer, 2021. doi:10.1007/978-3-030-79527-6_1.
- 4 Emmanuelle Anceaume, Yann Busnel, Ernst Schulte-Geers, and Bruno Sericola. Optimization results for a generalized coupon collector problem. *Journal of Applied Probability*, 53(2):622–629, 2016. doi:10.1017/jpr.2016.27.

- 5 Emmanuelle Anceaume, Yann Busnel, and Bruno Sericola. New results on a generalized coupon collector problem using markov chains. *Journal of Applied Probability*, 52(2):405–418, 2015. doi:10.1239/jap/1437658606.
- 6 Srinivasan Arunachalam, Aleksandrs Belovs, Andrew M Childs, Robin Kothari, Ansis Rosmanis, and Ronald de Wolf. Quantum coupon collector. In *15th Conference on the Theory of Quantum Computation, Communication and Cryptography*, 2020. doi:10.4230/LIPIcs.TQC.2020.10.
- 7 Siva Athreya, Satyaki Mukherjee, and Soumendu Sundar Mukherjee. On a generalisation of the coupon collector problem. *Journal of Theoretical Probability*, 38(3):54, 2025. doi:10.1007/s10959-025-01417-w.
- 8 Petra Berenbrink and Thomas Sauerwald. The weighted coupon collector's problem and applications. In *International Computing and Combinatorics Conference*, pages 449–458. Springer, 2009. doi:10.1007/978-3-642-02882-3_45.
- 9 Shoham Shimon Berrebi, Eitan Yaakobi, Zohar Yakhini, and Daniella Bar-Lev. The labeled coupon collector problem with random sample sizes and partial recovery. *arXiv preprint arXiv:2502.02968*, 2025. doi:10.48550/arXiv.2502.02968.
- 10 Arnon Boneh and Micha Hofri. The coupon-collector problem revisited—a survey of engineering problems and computational methods. *Stochastic Models*, 13(1):39–66, 1997. doi:10.1080/15326349708807412.
- 11 Peter J Brockwell. The extinction time of a birth, death and catastrophe process and of a related diffusion model. *Advances in Applied Probability*, 17(1):42–52, 1985. doi:10.2307/1427051.
- 12 Gregory Dobson and Tolga Tezcan. Optimal sampling strategies in the coupon collector's problem with unknown population size. *Annals of Operations Research*, 233(1):77–99, 2015. doi:10.1007/s10479-014-1563-0.
- 13 Devdatt P. Dubhashi and Alessandro Panconesi. *Concentration of measure for the analysis of randomized algorithms*. Cambridge University Press, 2009.
- 14 Pál Erdős and Alfréd Rényi. On a classical problem of probability theory. *A Magyar Tudományos Akadémia Matematikai Kutató Intézetének Közleményei*, 6(1-2):215–220, 1961. URL: <https://real.mtak.hu/id/eprint/201134>.
- 15 William Feller et al. *An introduction to probability theory and its applications*, volume 963. Wiley New York, 1971.
- 16 Marco Ferrante and Monica Saltalamacchia. The coupon collector's problem. *Materials mathematics*, pages 0001–35, 2014.
- 17 Philippe Flajolet, Danièle Gardy, and Loÿs Thimonier. Birthday paradox, coupon collectors, caching algorithms and self-organizing search. *Discrete Applied Mathematics*, 39(3):207–229, 1992. doi:10.1016/0166-218X(92)90177-C.
- 18 Gene H Golub and Charles F Van Loan. *Matrix computations*. JHU press, 2013.
- 19 Lars Holst. On birthday, collectors', occupancy and other classical urn problems. *International Statistical Review/Revue Internationale de Statistique*, pages 15–27, 1986. doi:10.2307/1403255.
- 20 Jelena Jocković and Bojana Todić. Coupon collector problem with reset button. *Mathematics*, 12(2):239, 2024. doi:10.3390/math12020239.
- 21 Tipaluck Krityakierne and Thotsaporn Aek Thanatipanonda. The slowest coupon collector's problem. *Journal of Systems Science and Complexity*, pages 1–13, 2025. doi:10.1007/s11424-025-4223-3.
- 22 David A Levin and Yuval Peres. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.
- 23 Stephen N Luko. The “coupon collector's problem” and quality control. *Quality Engineering*, 21(2):168–181, 2009. doi:10.1080/08982110802642555.
- 24 Michael Mitzenmacher and Eli Upfal. *Probability and computing: Randomization and probabilistic techniques in algorithms and data analysis*. Cambridge university press, 2017.
- 25 Donald J Newman. The double dixie cup problem. *The American Mathematical Monthly*, 67(1):58–61, 1960. URL: <http://www.jstor.org/stable/2308930>.

- 26 Prabhakar Raghavan and Rajeev Motwani. *Randomized algorithms*. Cambridge University Press Cambridge, 1995.
- 27 Ran Raz and Wei Zhan. The Random-Query Model and the Memory-Bounded Coupon Collector. In Thomas Vidick, editor, *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, volume 151 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 20:1–20:11, Dagstuhl, Germany, 2020. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ITCS.2020.20.
- 28 Bojana Todić. Coupon collector problem with penalty coupon. *Matematički vesnik*, 2023. doi:10.57016/MV-BGON6192.
- 29 Hermann Von Schelling. Coupon collecting for unequal probabilities. *The American Mathematical Monthly*, 61(5):306–311, 1954. doi:10.1080/00029890.1954.11988466.
- 30 Weiyu Xu and A Kevin Tang. A generalized coupon collector problem. *Journal of Applied Probability*, 48(4):1081–1094, 2011. doi:10.1239/jap/1324046020.