Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Section 1.
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes] All missing proofs can be found in the Appendix.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] The code used to run the experiments can be found in the supplemental material.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 6.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No] The experiments were simple enough that they were all completed within a few hours in our personal computers.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Omitted Proofs

A.1 Proof of Lemma 3.1

We split the statement into two separate parts.

Claim A.1.

$$\ln\left(1 - \frac{a_k}{2}\right) \geq -a_k$$

Proof. Consider the function $f : [0,1] \to \mathbb{R}_{\geq 0}$, where $f(x) = e^x\left(1 - x/2\right)$. Clearly, if $f(x) \geq 1$ for all $x \in [0,1]$, then the claim follows by taking the (natural) logarithm of each side of the inequality, and setting $x = a_k$.

We have $\frac{df(x)}{dx} = {-e^x}/{2}\,(x-1) \geq 0$ for all $x \in [0,1]$. Therefore, f is increasing in $[0,1]$, and thus attains its minimum for $x = 0$. Therefore, $f(x) \geq f(0) = 1$ for all $x \in [0,1]$ and the claim follows. $\square$

Claim A.2. *For every $j \in \{1, 2, \ldots, k-1\}$, we have*

$$\ln\left(1 - a_j + \frac{a_j^2}{2}\right) \geq -a_j$$

Proof. Fix an arbitrary a_j. Consider the function $g : [0,1] \to \mathbb{R}_{\geq 0}$, where $g(x) = e^x\left(1 - x + {x^2}/{2}\right)$. Clearly, if $g(x) \geq 1$ for all $x \in [0,1]$, then the claim follows by taking the (natural) logarithm of each side of the inequality, and setting $x = a_j$.

We have $\frac{dg(x)}{dx} = {e^x x^2}/{2} \geq 0$ for all $x \in [0,1]$. Therefore, g is increasing in $[0,1]$, and thus attains its minimum for $x = 0$. Therefore, $g(x) \geq g(0) = 1$ for all $x \in [0,1]$ and the claim follows. $\square$

A.2 Proof of Lemma 3.2

π is clearly a randomized OCRS because every time it sees an element, it makes an irrevocable decision to select it, if it is active, before it sees the next element, and also, by the choice of $\mathcal{F}_{\pi,x}$, it is easy to see that the set of elements it returns is always a singleton, and thus feasible in $\mathcal{I}$, since $\mathcal{F}_{\pi,x} \subseteq \mathcal{I}$. Furthermore, the choice of $\mathcal{F}_{\pi,x}$ is randomized, and thus π is a randomized OCRS.

Next, it is also easy to see that π is a greedy OCRS, because, given x, $\mathcal{F}_{\pi,x}$ is a down-closed subfamily of feasible sets and an active element e is always selected if $\{e\} \in \mathcal{F}_{\pi,x}$, since there are no previously selected elements.

A.3 An alternate proof of Theorem 1.1

The following scheme is due to Jan Vondrák [Von].

Let π denote the OCRS we will create. π will draw a random set R where each element e_i appears in R independently with some probability q_i. Afterwards, it will set

$$\mathcal{F}_{\pi,x} = \{\{e_i\} \mid e_i \in R\}\,.$$

We set $q_i = {1-e^{-x_i}}/{x_i}$ for all $e_i \in \mathcal{N}$. Afterwards, π selects the first element e_i that is active and that $\{e_i\} \in \mathcal{F}$.

The proof of the next lemma is identical to the proof of Lemma 3.2.

Lemma A.3. *π is a randomized greedy OCRS.*

Next, we quantify the probability that each element is selected by π, given that it is active.

Lemma A.4. *π selects every element $e_i \in \mathcal{N}$, given that it is active, with probability at least ${1}/{e}$.*

Proof. We relabel the elements of $\mathcal{N}$ so that each e_i arrives in the i-th step. Consider an element $e_i \in \mathcal{N}$. Given that e_i is active, since π is a greedy OCRS, π will select e_i if and only if it has not selected any elements before e_i and also $\{e_i\} \in \mathcal{F}_{\pi,x}$. Recall that we have $\{e_i\} \in \mathcal{F}_{\pi,x}$ with probability exactly $q_i = {1-e^{-x_i}}/{x_i}$. Furthermore, for every element e_j where $j < i$, it needs to be the case that we avoid having both $\{e_j\} \in \mathcal{F}_{\pi,x}$ and also e_j coming up active. This happens with probability $1 - x_j \cdot {1-e^{-x_j}}/{x_j} = e^{-x_j}$ for every e_j where $j < i$. Overall, if we denote by r_i the probability that e_i is selected by π, given that it is active, we have

$$r_i = \frac{1-e^{-x_i}}{x_i} \cdot \prod_{j<i} e^{-x_j} = \frac{1-e^{-x_i}}{x_i} \cdot e^{-\sum_{j<i} x_j} \geq \frac{(1-e^{-x_i})\,e^{x_i - 1}}{x_i} = \frac{e^{x_i-1} - e^{-1}}{x_i},$$

where the inequality follows from $\sum_i x_i \leq 1$. This expression is minimized for $x_i \to 0$, and thus we get $r_i \geq {1}/{e}$, for all $i \in \mathcal{N}$. $\square$

From Lemmas A.3 and A.4, it follows that π is a ${1}/{e}$-selectable (randomized) greedy OCRS for $\mathcal{P}$.

Remark A.5. One can easily see that the difference between the two proofs is that, in our scheme, the probability of selection q_i of each element $i \in \mathcal{N}$ is a linear approximation of the selection probability of Vondrák's scheme. The result then follows due to the convexity of the selection probability $q_i = {}^{1-e^{-x_i}}/_{x_i}$ of Vondrák's scheme.

A.4 Proof of Lemma 4.1

Assume towards contradiction, that

$$\min_{e \in \mathcal{N}} \left\{ \sum_{k=1}^{n} \left(1 - \frac{1}{n}\right)^{k-1} \sum_{\substack{S \subseteq \mathcal{N} \,:\, |S|=k \\ e \in S}} \alpha_S \right\} > \left(1 - \frac{1}{n}\right)^{n-1}.$$

The proof consists of a double counting argument. First, notice that, by the inequality above, we have

$$\sum_{e \in \mathcal{N}} \left(\sum_{k=1}^{n} \left(1 - \frac{1}{n}\right)^{k-1} \sum_{\substack{S \subseteq \mathcal{N} \,:\, |S|=k \\ e \in S}} \alpha_S \right) > n \left(1 - \frac{1}{n}\right)^{n-1}. \tag{1}$$

For any $0 \leq k \leq n$, let $\beta_k = \sum_{S \subseteq \mathcal{N} \,:\, |S|=k} \alpha_S$ be the total probability mass assigned by the greedy OCRS to all sets of size k, and notice that $\sum_{k=0}^{n} \beta_k = 1$. We can also compute the left-hand side of (1) as

$$\begin{aligned}
\sum_{e \in \mathcal{N}} \left(\sum_{k=1}^{n} \left(1 - \frac{1}{n}\right)^{k-1} \sum_{\substack{S \subseteq \mathcal{N} \,:\, |S|=k \\ e \in S}} \alpha_S \right) &= \sum_{k=1}^{n} \left(\left(1 - \frac{1}{n}\right)^{k-1} \sum_{e \in \mathcal{N}} \sum_{\substack{S \subseteq \mathcal{N} \,:\, |S|=k \\ e \in S}} \alpha_S \right) \\
&= \sum_{k=1}^{n} \left(k \left(1 - \frac{1}{n}\right)^{k-1} \sum_{S \subseteq \mathcal{N} \,:\, |S|=k} \alpha_S \right) \\
&= \sum_{k=1}^{n} \left(\beta_k \cdot k \left(1 - \frac{1}{n}\right)^{k-1} \right).
\end{aligned} \tag{2}$$

where the second equality follows from the fact that in the double sum, for every S such that $|S| = k$, every coefficient a_S appears exactly k times, one for each element it contains. Under the constraint $\sum_{k=0}^{n} \beta_k = 1$, we have that $\sum_{k=1}^{n} \left(\beta_k \cdot k \left(1 - {}^{1}/_{n}\right)^{k-1} \right)$ is maximized for $\beta_n = 1$ and $\beta_m = 0$ for all $m < n$, as $k \left(1 - {}^{1}/_{n}\right)^{k-1}$ is strictly increasing in k. Therefore,

$$\sum_{e \in \mathcal{N}} \left(\sum_{k=1}^{n} \left(1 - \frac{1}{n}\right)^{k-1} \sum_{\substack{S \subseteq \mathcal{N} \,:\, |S|=k \\ e \in S}} \alpha_S \right) \leq n \left(1 - \frac{1}{n}\right)^{n-1}. \tag{3}$$

Together, (1) and (3) yield a contradiction.

A.5 Proof of Lemma 5.2

Consider an active element $u \in U$. Since π is a greedy OCRS, it will select u if and only if there exists a neighbor v of u such that $u \in R_v$, and also, together with the set S of elements already selected by π, $S + u \in \mathcal{F}$. First, for every element $w \in U$, let $\mathcal{E}_w$ denote the event that there exists an element $v \in V$ such that $w \in R_v$. In other words, $\mathcal{E}_w$ is the event that w is in some set of $\mathcal{F}$. For $\mathcal{E}_u$, we have

$$\Pr[\mathcal{E}_u] = 1 - \prod_{v \in N(u)} (1 - q_u) = 1 - (1 - q_u)^{|N(u)|} = \frac{1 - e^{-x_u}}{x_u}.$$

Furthermore, the set S selected prior to seeing u has to be independent, thus $S \in \mathcal{F}$, and thus for $S + u \notin \mathcal{F}$, it has to be that for every $v \in N(u)$, we have $|S \cap R_v| \geq 1$. Therefore, the probability that $S + u \notin \mathcal{F}$ is

$$\begin{aligned}
\Pr\left[S + u \notin \mathcal{F} \mid S\right] &= \prod_{v \in N(u)} \left(1 - \prod_{\substack{u' \in N(v) \\ u' \neq u}} \left(1 - x'_u \Pr[\mathcal{E}_{u'}]\right)\right) \\
&= \prod_{v \in N(u)} \left(1 - \prod_{\substack{u' \in N(v) \\ u' \neq u}} \left(1 - x'_u \frac{1 - e^{-x_{u'}}}{x_{u'}}\right)\right) \\
&= \prod_{v \in N(u)} \left(1 - \prod_{\substack{u' \in N(v) \\ u' \neq u}} e^{-x_{u'}}\right) \\
&= \prod_{v \in N(u)} \left(1 - e^{-\sum_{u' \in N(v): u' \neq u} x_{u'}}\right) \\
&\leq \prod_{v \in N(u)} \left(1 - e^{-1 + x_u}\right) \\
&= \left(1 - e^{-1 + x_u}\right)^{|N(u)|},
\end{aligned}$$

where the inequality follows from the fact that, for every $v \in V$, $\sum_{w \in N(v)} x_w \leq 1$ due to $x \in \mathcal{P}$. Therefore, we have

$$\begin{aligned}
\Pr[u \in \pi(R) | u \in R] &= \Pr[\mathcal{E}_u] \cdot \Pr\left[S + u \in \mathcal{F} \mid S\right] \\
&\geq \frac{1 - e^{-x_u}}{x_u} \left(1 - \left(1 - e^{-1 + x_u}\right)^{|N(u)|}\right).
\end{aligned}$$

Let $f_k(x) = {}^{1 - e^{-x}}\!/_{x} \left(1 - \left(1 - e^{-1+x}\right)^k\right)$, for $k \geq 1$ and $x \in [0, 1]$. It is easy to see that $f_k(x) \geq {}^1\!/_e$ for every $k \geq 1$ and $x \in [0, 1]$. Furthermore, we have that for $k \geq 3$, $f_k(x)$ is minimized in $[0, 1]$ for $x = 1$, and yields $f_k(1) = 1 - {}^1\!/_e$.