## A Proof of Theorem 1.2

LP Construction First, consider the following LP construction for the learning-augmented skirental problem: We use $\gamma$ to denote the robustness parameter and $\beta$ the consistency parameter. We assume without loss of generality that $\beta<\gamma$; otherwise, the consistency requirement is redundant. Consider a infinite LP, with variables $\left\{p_{i}\right\}$ indicating the probability of buying at day $i$. First, we ask that the $p_{i}$ 's define a probability distribution. That is, $p_{i} \geq 0$ and

$$
\begin{equation*}
\sum_{i=1}^{\infty} p_{i}=1 \tag{probabilitydistribution}
\end{equation*}
$$

Second, to satisfy the consistency constraint, the algorithm must have expected cost within $\beta$. OPT when $y=x$. In this case, the ski season ends at $i=y$, so there is no additional cost afterwards.

$$
\begin{equation*}
\sum_{i=1}^{y}(B+i-1) p_{i}+y \sum_{i=y+1}^{\infty} p_{i} \leq \beta \min \{B, y\} \tag{consistency}
\end{equation*}
$$

Third, each value of $x$ gives a distinct constraint for robustness, where the left side is the expected cost and the right side is $\gamma$. OPT. When $x \leq B$, OPT $=x$, so we have

$$
\begin{equation*}
\sum_{i=1}^{x}(B+i-1) p_{i}+x \sum_{i=x+1}^{\infty} p_{i} \leq \gamma x \quad \forall x \leq B \tag{A.1}
\end{equation*}
$$

If $x>B$, then OPT $=B$. The robustness constraints are infinitely many, given by

$$
\begin{equation*}
\sum_{i=1}^{x}(B+i-1) p_{i}+x \sum_{i=x+1}^{\infty} p_{i} \leq \gamma B \quad \forall x>B \tag{A.2}
\end{equation*}
$$

We remark that for each $p_{i}$, its coefficient is non-decreasing as we go down. We denote the robustness constraint corresponding to $x$ by $\mathcal{C}(x)$ and the entire (infinite) LP by $\mathcal{P}$.
From now on, we focus on the case when $y \geq 2 B-1$. In this case, the consistency constraint is

$$
\sum_{i=1}^{y}(B+i-1) p_{i}+y \sum_{i=y+1}^{\infty} p_{i} \leq \beta B
$$

Reducing to a Finite LP We will show that $\mathcal{P}$ can be reduced to be a finite LP of $B+1$ constraints and $y$ variables.
Lemma A.1. If $y \geq B$, the robustness constraints between $\mathcal{C}(B)$ and $\mathcal{C}(y)$ are redundant.

Proof. Observe that these constraints are dominated by the constraint consistency', since $\beta \leq \gamma$ and their left sides are bounded by the left side of consistency'.

We start by identifying redundant variables in $\mathcal{P}$.
Lemma A. 2 (Redundant variables). If $y \geq 2 B-1$ and $\mathcal{P}$ is feasible, then there exists a feasible solution to $\mathcal{P}$ such that $p_{i}=0$ for all $i \geq y+1$.

Proof. Let $p$ be a feasible solution to $\mathcal{P}$. First, we argue that there exists a feasible solution $p^{\prime}$ such that $p_{y+2}^{\prime}=p_{y+3}^{\prime}=\cdots=0$. To eliminate $p_{y+2}$, consider $p^{\prime}$ where $p_{y+2}^{\prime}=0, p_{y+1}^{\prime}=p_{y+1}+p_{y+2}$, and $p_{i}^{\prime}=p_{i}$ for $i \notin\{y+1, y+2\}$. Clearly, $p^{\prime}$ still defines a probability distribution. We now check $p^{\prime}$ is feasible.
(i) Since $y \geq B$, the consistency constraint is satisfied.
(ii) The robustness constraints from $\mathcal{C}(1)$ to $\mathcal{C}(B-1)$ are satisfied, by the coefficients in these constraints.
(iii) By Lemma A.1, we focus on the robustness constraints from $\mathcal{C}(y+1)$. First, the coefficient of $p_{y+2}$ is greater than the coefficient of $p_{y+1}$ in the constraint $\mathcal{C}(i)$. for all $i \geq y+2$. Hence, $p^{\prime}$ satisfies these constraints. Then, note that the constraint $\mathcal{C}(y+1)$ is dominated by the constraint $\mathcal{C}(y+2)$, and thus $p^{\prime}$ satisfies it.

Applying this argument iteratively, we can eliminate all variables $p_{i}$ for $i \geq y+2$. Now it is easy to observe that constraints $\mathcal{C}(j)$ for $j \geq y+2$ are redundant.
Finally, to eliminate $p_{y+1}$, consider $p_{y+1}^{\prime}=0, p_{y-B+1}^{\prime}=p_{y-B+1}+p_{y+1}$, and $p_{i}^{\prime}=p_{i}$ for $i \notin$ $\{y-B+1, y+1\}$. Observe that since $y \geq 2 B-1$,
(i) in the consistency' constraint and constraint $\mathcal{C}(y+1)$, the coefficient of $p_{y+1}$ and $p_{y-B+1}$ are both $y$;
(ii) in the constraints between $\mathcal{C}(1)$ and $\mathcal{C}(B-1)$, the coefficient of $p_{y+1}$ and $p_{y-B+1}$ are both $B$.

It follows that all constraints are satisfied by $p^{\prime}$.
Corollary A. 3 (Redundant constraints). If $y \geq 2 B-1$, all robustness constraints in $\mathcal{P}$ are redundant except those between $\mathcal{C}(1)$ and $\mathcal{C}(B-1)$.

Proof. This follows directly from the definition of the robustness constraints and Lemma A.2.
Lemma A. 4 (Tight constraints). If $\mathcal{P}$ is feasible, there exists a solution such that for all $1<k<y$, if $p_{k}>0$, then the preceding robustness constraints $\mathcal{C}\left(k^{\prime}\right)$ for $1 \leq k^{\prime}<k$ are all tight.

Proof. Using the probability distribution constraint, we can rewrite each robustness constraint $\mathcal{C}(i)$ as

$$
\begin{equation*}
(B-i) p_{1}+(B-i+1) p_{2}+\cdots+(B-1) p_{i} \leq \gamma \min \{B, i\}-i \tag{A.3}
\end{equation*}
$$

Let $p$ be a feasible solution. First, we claim that when shifting probability mass from $p_{k}$ to $p_{k-1}$, the slack for all robustness constraints is non-decreasing, except for $\mathcal{C}(k-1)$. Note that
(i) for $k^{\prime} \geq k$, since the coefficient of $p_{k-1}$ in $\mathcal{C}\left(k^{\prime}\right)$ is less than that of $p_{k}$, we strictly increase the slack; and
(ii) for $k^{\prime}<k-1, \mathcal{C}\left(k^{\prime}\right)$ has no dependence on either $p_{k}$ or $p_{k-1}$, so the slack remains unchanged.

Second, we claim that if $\mathcal{C}(k-1)$ has non-zero slack, then $p_{k}>0$ or $\mathcal{C}(k)$ has non-zero slack. Indeed, if $p_{k}=0$, then constraint $\mathcal{C}(k-1)$ is stronger than constraint $\mathcal{C}(k)$. Since constraint $\mathcal{C}(k-1)$ has non-zero slack, constraint $\mathcal{C}(k)$ must also have non-zero slack.
Let $s=\left(s_{1}, s_{2}, \ldots, s_{y-1}\right)$ be the vector of slacks for the robustness constraints $\mathcal{C}(i)$. The above two claims together show that if a feasible solution is such that $p_{k}>0$, but one of the preceding robustness constraints is not tight, then we can shift some probability mass so that the robustness slack vector becomes lexicographically smaller. (Here, we consider the lexicographic ordering on $\mathbb{R}^{y-1}$.) Equivalently, if the robustness slack vector is lexicographically minimal, then $p_{k}>0$ implies all of the preceding robustness constraints are tight. It thus suffices to show that a lexicographically minimal robustness slack vector exists.
Let $S$ be the set of robustness slack vectors that correspond to feasible solutions. Observe that $S$ is compact. The existence of a lexicographically minimal element in $S$ then follows because any compact subset of $\mathbb{R}^{y-1}$ contains a lexicographically minimal element. To see this last point, note that compactness implies there exists a element with minimum first coordinate. Now, restrict our compact set to this minimum first coordinate and repeat this argument for the second coordinate, and so on.

Finally, to prove our main theorem, we need the following technical lemma.
Lemma A.5. For all $B>1$ and $x \in[0,1]$, the following inequality holds:

$$
\frac{1}{B} x-\left(1+\frac{1}{B-1}\right)^{-1}\left(\left(1+\frac{1}{B-1}\right)^{x}-1\right) \geq 0
$$

Proof. The stated inequality is equivalent to

$$
\left(1+\frac{1}{B-1}\right) \frac{1}{B} x+1 \geq\left(1+\frac{1}{B-1}\right)^{x}
$$

Observe that the left-hand side is linear and that the right-hand side is convex. Since the two sides are equal at $x=0$ and $x=1$, the desired inequality is true on the interval $[0,1]$ by Jensen's inequality.

Now we are ready to present the proof of the randomized lower bound.

Proof of Theorem 1.2. Fix a cost $B$. Given any value of robustness $\gamma>1$, we give a lower bound for the consistency $\beta$ in the "hard" case $y=2 B-1$. Assume for now that the LP with parameters $\gamma$ and $\beta$ is feasible. We will derive constraints on $\beta$ in terms of $\gamma$.

By Lemma A.2, we know that there exists a feasible solution $p$ with $p_{y+1}=0$. By Lemma A.4, there exists a $k \leq y$ such that $p_{i}>0$ for $1 \leq i \leq k$ and $p_{i}=0$ for all $i>k$. Moreover, the first $k-1$ constraints of the LP all have 0 slack. In fact, $k \leq B$ always: By our feasibility assumption, $\gamma$ must be at least the optimal competitive ratio $c^{*}=e /(e-1)$ in the classic setting, since that setting has no consistency constraint. In the classic setting, we can achieve the optimal competitiveness with $k=B$ and have the first $k$ robustness constraints be tight [KMMO94]. Thus, with $\gamma \geq c^{*}$, the robustness constraints are relaxed, so we must also have $k \leq B$. With these observations and the constraint $p_{1}+p_{2}+\cdots+p_{k}=1$, we can determine the value of $k$ and the probabilities $p_{1}, \ldots, p_{k}$.
It is not difficult to see via induction on $k$ (e.g., using the reformulation of Equation (A.3)) that for the first $k-1$ constraints to each have 0 slack, we must have

$$
p_{i}=\frac{\gamma-1}{B-1}\left(1+\frac{1}{B-1}\right)^{i-1}
$$

for each $i$ between 1 and $k-1$. From $p_{1}+p_{2}+\cdots+p_{k}=1$, it follows that $k$ must be the smallest integer such that

$$
\sum_{i=1}^{k} \frac{\gamma-1}{B-1}\left(1+\frac{1}{B-1}\right)^{i-1}=(\gamma-1)\left(\left(1+\frac{1}{B-1}\right)^{k}-1\right) \geq 1
$$

Rearranging the inequality now gives us

$$
k=\left\lceil\frac{\log \left(1+\frac{1}{\gamma-1}\right)}{\log \left(1+\frac{1}{B-1}\right)}\right\rceil \text {. }
$$

This choice of $k$, our definition of $p_{i}$ for $1 \leq i \leq k-1$, and the constraint $\sum_{i=1}^{k} p_{i}=1$ fully determine a feasible solution.
By our assumption that the LP is feasible, this setting of $p_{i}$ values must also satisfy the consistency constraint. That is,

$$
\begin{equation*}
\sum_{i=1}^{k}(B+i-1) p_{i}=B+\sum_{i=1}^{k}(i-1) p_{i} \leq \beta B \tag{A.4}
\end{equation*}
$$

The remainder of this proof consists of computing the left-hand side of the above for our feasible solution explicitly to obtain a lower bound for the consistency $\beta$.

Applying our explicit formulas for $p_{i}$ for $1 \leq i \leq k-1$ and the fact $p_{k}=1-\sum_{i=1}^{k-1} p_{i}$, we compute the sum

$$
\begin{aligned}
\sum_{i=1}^{k}(i-1) p_{i}= & \left(\sum_{i=1}^{k-1}(i-1) p_{i}\right)+(k-1)\left(1-\sum_{i=1}^{k-1} p_{i}\right) \\
= & (\gamma-1)\left(B-(B-k+1)\left(1+\frac{1}{B-1}\right)^{k-1}\right) \\
& \quad+(k-1)\left(1-(\gamma-1)\left(\left(1+\frac{1}{B-1}\right)^{k-1}-1\right)\right) \\
= & (k-1) \gamma+(\gamma-1) B\left(1-\left(1+\frac{1}{B-1}\right)^{k-1}\right)
\end{aligned}
$$

It follows from Equation (A.4) that $\beta$ is lower bounded by

$$
\begin{aligned}
\beta & \geq 1+\frac{1}{B} \sum_{i=1}^{k}(i-1) p_{i} \\
& =1+\frac{(k-1) \gamma}{B}+(\gamma-1)\left(1-\left(1+\frac{1}{B-1}\right)^{k-1}\right)
\end{aligned}
$$

Now, define

$$
\Delta k:=k-\frac{\log \left(1+\frac{1}{\gamma-1}\right)}{\log \left(1+\frac{1}{B-1}\right)} .
$$

After some further computation, we obtain

$$
\begin{aligned}
\beta & \geq \frac{k \gamma}{B}-\gamma\left(1-\frac{1}{B}\right)\left(\left(1+\frac{1}{B-1}\right)^{\Delta k}-1\right) \\
& =\frac{\gamma}{B} \frac{\log \left(1+\frac{1}{\gamma-1}\right)}{\log \left(1+\frac{1}{B-1}\right)}+\gamma\left(\frac{\Delta k}{B}-\left(1+\frac{1}{B-1}\right)^{-1}\left(\left(1+\frac{1}{B-1}\right)^{\Delta k}-1\right)\right)
\end{aligned}
$$

Lemma A. 5 lets us bound the terms involving $\Delta k$ from below, and it follows that

$$
\beta \geq \frac{\gamma}{B} \frac{\log \left(1+\frac{1}{\gamma-1}\right)}{\log \left(1+\frac{1}{B-1}\right)}
$$

To finish our proof of the lower bound, notice that $B \log (1+1 /(B-1)) \rightarrow 1$ from below as $B \rightarrow \infty$. Hence, the lower bound on $\beta$ approaches $\gamma \log (1+1 /(\gamma-1))$ as $B \rightarrow \infty$.

## B Proof of Theorem 3.2

Now we present our algorithmic result. Although our analysis deals with the case of 2 jobs, it is convenient to describe the algorithm in the general case of $n$ jobs. The algorithm starts by running round robin for a while, then switches to a greedy strategy of processing jobs in the increasing order of the predicted times. If at any point we know $x_{i} \neq y_{i}$ for any job $i$, we switch to round robin forever. We use $\mathrm{OPT}_{y}=\sum_{i} i y_{i}$ to denote the OPT under perfect predictions.

```
Two-Stage-Schedule}(\mp@subsup{y}{1}{},\mp@subsup{y}{2}{},\cdots,\mp@subsup{y}{n}{})\mathrm{ :
    At any point, if a job finishes with processing time less or more than its prediction,
        round robin forever.
    Stage 1: Round robin for at most \lambdan\cdot\mp@subsup{\textrm{OPT}}{y}{}/(\begin{array}{l}{n}\\{2}\end{array})\mathrm{ units of time.}
    Stage 2: Process jobs in predicted order
            (staring from the unfinished job with the least predicted time).
```

The intuition behind the algorithm is simple. On one hand, to ensure robustness, the algorithm switches to round robin when any misprediction is noticed. On the other hand, we ask the algorithm to be $(1+\lambda)$-consistent. Suppose $y_{1}<y_{2}<\cdots<y_{n}$. If the predictions are perfect, then we expect that a consistent algorithm would produce a schedule that finishes the jobs in the correct order, i.e., job 1 finishes first, job 2 second, and so on. In this case, the consistency requirement reduces to

$$
\begin{equation*}
\sum_{i>j} d(i, j) \leq \lambda \mathrm{OPT}_{y} \tag{B.1}
\end{equation*}
$$

where and $d(i, j)$ denotes the amount job $i$ delays job $j$ in this scenario. Observe that when no job is completed, round robin increases each term in the summation at the same rate of $1 / n$. Thus, stage 1 of the algorithm would make the inequality (B.1) tight. Then as we can no longer disobey the predictions in the ideal scenario, we switch to the greedy strategy in the second stage. Next, we analyze the performance of the algorithm in the case of two jobs.
We now prove Theorem 3.2
Proof of Theorem 3.2. Let $t=2 y_{1}+y_{2}$. To show consistency, assume $x_{1}=y_{1}, x_{2}=y_{2}$, so $\mathrm{OPT}=t$. In stage 1 , the algorithm runs round robin for $2 \lambda t$ units of time. Observe that job 2 cannot finish before job 1 in this stage: since $\lambda<1 / 3$, job 2 can receive at most $\left(2 y_{1}+y_{2}\right) / 3<y_{2}$ units of processing time. Consider two cases.
(i) Suppose job 1 finishes in stage 1. Then since two jobs share the same rate,

$$
\begin{equation*}
y_{1} \leq \lambda t . \tag{B.2}
\end{equation*}
$$

Moreover, in this case. the algorithm runs round robin for $2 y_{1}$ time and finishes job 2 in $y_{2}-y_{1}$ time. Thus, $\mathrm{ALG}=3 y_{1}+y_{2}$, and $\mathrm{OPT}=t$. By (B.2), we have $\mathrm{ALG} \leq(1+\lambda) \mathrm{OPT}$.
(ii) Suppose job 1 does not finish in stage 1 . Then both jobs have been processed for $\lambda t$ units of time at the beginning of stage 2 . In stage 2 , the algorithm prioritizes job 1 . Thus,

$$
\begin{equation*}
\mathrm{ALG}=4 \lambda t+2\left(y_{1}-\lambda t\right)+\left(y_{2}-\lambda t\right)=(1+\lambda) \mathrm{OPT} \tag{B.3}
\end{equation*}
$$

To show robustness, we consider mispredictions, and suppose without loss of generality $y_{1}=1$. Throughout, we let $\epsilon$ to denote an infinitesimal quantity. Notice that if any misprediction is found or job 1 is finished in stage 1 , the algorithm is equivalent of round robin and, therefore, achieves $4 / 3$ competitive ratio that is better than $1+1 /(1+6 \lambda)$ for any $\lambda \in(0,1 / 3)$, so we are done. We do a case-by-case analysis, assuming in stage 1 no misprediction is detected and both jobs are finished in stage 2 . Notice that under the assumptions, $x_{1}, x_{2} \geq \lambda t$, so OPT $\geq 3 \lambda t$.
(i) Suppose job 1 finishes no later than its prediction ( $x_{1} \leq 1$ ). We have $\operatorname{ALG}=\lambda t+2 x_{1}+x_{2}$.
(a) If $x_{1}<x_{2}$, then OPT $=2 x_{1}+x_{2}$. Since $\lambda t \leq$ OPT $/ 3$, we have $\mathrm{ALG} / \mathrm{OPT} \leq 4 / 3$.
(b) If $x_{1} \geq x_{2}$, then OPT $=2 x_{2}+x_{1}$. Observe that setting $x_{1}=y_{1}=y_{2}=1, x_{2}=\lambda t+\epsilon$ maximizes the competitive ratio, and this yields a ratio of $1+1 /(1+6 \lambda)$.
(ii) Suppose job 1 finishes later than its prediction $\left(x_{1}>1\right)$. In this case, the stage 2 starts off by processing job 1 for $y_{1}-\lambda t$ unit of time then switching to round robin.
(a) If job 1 finishes no later than job 2, then we calculate that $\mathrm{ALG}=\lambda t+3 x_{1}+x_{2}-1$. If $x_{1}<x_{2}$, then OPT $=2 x_{1}+x_{2}$, the competitive ratio is at most $4 / 3$, where the worst case is achieved at $x_{1}=1+\epsilon$ and we use $\lambda t \leq \mathrm{OPT} / 3$. If $x_{1} \geq x_{2}$, then OPT $=2 x_{2}+x_{1}$. The competitive ratio is bounded by $1+1 /(1+6 \lambda)$, where the worst case is achieved when $x_{1}=1+\epsilon, x_{2}=\lambda t+2 \epsilon, y_{2}=1$.
(b) If job 1 finishes later than job 2, then $\mathrm{ALG}=1+x_{1}+3 x_{2}-\lambda t$. Observe that in this case, it is impossible that $x_{2}>x_{1}$, since job 1 receives more processing than job 2 throughout the schedule. Assume $x_{2} \leq x_{1}$; then the competitive ratio is bounded by $1+1 /(1+6 \lambda)$ with the worst case being $x_{2}=\lambda t+\epsilon, x_{1}=1$.

