AMSI VACATIONRESEARCH SCHOLARSHIPS 2022–23

SET YOUR SIGHTS ON RESEARCH THIS SUMMER

Prophet Inequalities

Joel Denning

Supervised by A/Prof Georgy Sofronov & Mr. Hugh Entwistle

Macquarie University

AMSI AUSTRALIAN MATHEMATICAL SCIENCES INSTITUTE

Contents

1 Prelude

Acknowledgements

I would like to thank Hugh Entwistle for his continued support in developing my understanding of the topics studied, especially when I would continually misunderstand a concept and ask mundane questions; his patience and explanations were appreciated.

Also, to Georgy Sofronov, thank you for engaging me in probability and statistics over the course of 2022 through my university studies, I would never have thought that I would enjoy these branches of mathematics as much as I do without you teaching me.

Finally, thank you to my family for pretending to understand what I'm talking about and making silly jokes any time probability is brought up in conversation, and thank you to my girlfriend, Jemma, for your amazing support and encouragement in taking on this project.

Abstract

We develop an understanding of the theory of optimal stopping so that prophet inequalities, a comparison between the expected maximum of a sequence of random variables and the reward gained using a non-anticipating stopping rule, may be studied. In doing so, we also discuss order statistics and stopping rules before proving numerous important results and demonstrating them empirically.

Statement of Authorship

All results and information in this report can be found in various texts referenced at the end of this report. Notations, structure, and language have been updated to reflect a simpler and more consistent approach to the results.

All updates to notation and elaborations in proofs were implemented by Mr. Denning, with assistance from Mr. Entwistle and Dr. Sofronov in understanding the original language used and decomposing the proofs into lighter strands. Any code used was developed by Mr. Denning with reference to previous work undertaken by Mr. Entwistle.

2 Introduction

Many problems exist in finance, economics, and other applications, where data is collected and analysed sequentially. This data is collected in real-time and we may require a decision to be made before we have had a chance to look at all the data; a decision must be made without seeing future observations and are based only on past and present observations, and a choice of returning to a previous point in the sequence may not be allowed. These problems may not have an obvious solution; we may instead employ a stopping rule, where we analyse the collected data until a set of criteria have been satisfied, at which point we stop and cease further action (hence accepting some reward or reducing cost of further research). We would then be concerned with how good our decision was relative to the best possible outcome.

Prophet inequalities are a class of problems in the theory of optimal stopping which allow us to answer the above question. The idea for this report is that we are given n random variables $X_1, \ldots, X_n$ with known distributions but unknown realisations $x_1, \ldots, x_n$. We observe the values one-by-one and we must choose a time at which to cease observation and accept the value we are on, with the goal of maximising this value. To do so, we employ a strategy (via a stopping rule) that decides to either choose a value x_i $(i = 1, \ldots, n)$ knowing only that and any previously observed values, or to reject x_i and move on to the next value, at which point we cannot return to x_i (or any other previously observed values) again. If we then consider a game of this nature through the lens of a gambler (someone using an optimal non-anticipating stop rule) versus a prophet (someone with complete foresight who can always win), then prophet inequalities allow us to compare the average performance of this gambler against the prophet.

In this paper, we look at and prove some well-known results before experimentally demonstrating and verifying the inequalities. As these inequalities were first proven many years ago, we aim to advance some of the terminology and notation to the present and provide more approachable proofs of these results whilst relating them to specific examples. A nice summary of some general prophet results may be found in [5], and an introduction to optimal stopping in [1].

3 Background

3.1 Optimal Stopping and Stopping Rules

The theory of optimal stopping is concerned with choosing a time to take an action to maximise an expected reward or minimise an expected cost. This action is based upon sequentially observing a number of random variables. In close relation to optimal stopping, stopping rules are then a set of criteria that specify whether we should cease observation and choose the current value, or to continue on to the next observation.

To determine when to stop, we will observe the sequence of random variables and decide, at each step, whether to take x_i, the current value of X_i, or continue on. The problem of finding the optimal stopping rule can be solved via backward induction using the following recurrence relation:

$$v_i = E(\max\{X_{n-i+1}, v_{i-1}\}), \quad i = 1, \ldots, n, \;\; v_0 = -\infty,$$

where v_i is the expected reward of a sequence with i observations left to make, and v_n is the overall expected reward. Intuitively, with i observations left to make, we could either accept the next value, x_{n-i+1}, or reject it and continue on knowing the expected reward with $i - 1$ observations remaining, and we should choose the maximum of those.

Before continuing, it will be helpful to give the above expression for v_i differently. For $X_1, \dots, X_n$ that are independent and identically distributed (iid) and continuous random variables, then by the definition of expected value,

$$\begin{aligned} v_i &= E(\max\{X_{n-i+1}, v_{i-1}\}) \\ &= \int_{-\infty}^{\infty} \max\{x, v_{i-1}\} f(x)\, dx, \end{aligned}$$

where $f(x)$ is the common probability density function (pdf) of each X_i.

The corresponding stopping rule which allows us to obtain the maximum expected reward with stopping value $\tau^* \in T_n$, where T_n is the set of all stopping rules for $X_1, \dots, X_n$, is

$$\tau^* = \min\{i \; : \; 1 \le i \le n,\, x_i \ge v_{n-i}\}.$$

What this rule does is stops at the first time the condition $x_i \ge v_{n-i}$ is met; we compare the current value x_i with the expected reward if we were to reject that value and move on, v_{n-i}. It may be shown that this rule is the optimal rule under our assumptions of independence [7].

Example 3.1. Let $X_1, \dots, X_n$ be iid uniform random variables on [a,b]. These have common pdf $f(x) = (b-a)^{-1}$ if $x \in [a, b]$, and $= 0$ elsewhere. We will now compute the v_i values using the above relation:

$$\begin{aligned} v_i &= E(\max\{X_{n-i+1}, v_{i-1}\}) \\ &= \int_a^b (b-a)^{-1} \max\{x, v_{i-1}\}\, dx \\ &= (b-a)^{-1} \left\{ \int_a^{v_{i-1}} v_{i-1}\, dx + \int_{v_{i-1}}^b x\, dx \right\} \\ &= \frac{(v_{i-1} - a)^2}{2(b-a)} + \frac{a+b}{2} \end{aligned}$$

for $i = 1, \dots n$.

For the sake of concreteness, consider the [0,1] uniform distribution (so $a = 0,\ b = 1$ and hence $v_i = \frac{(v_{i-1})^2+1}{2}$) and let $n = 7$. This gives us our values

n	0	1	2	3	4	5	6	7
v_i	0.0000	0.5000	0.6250	0.6953	0.7417	0.7751	0.8004	0.8203

where $v_7 = 0.8203$ is our expected reward. We hence have the optimal stopping rule τ^*, where

$$\tau^* = \min\{i \; : \; 1 \le i \le 7,\ x_i \ge v_{7-i}\}.$$

For instance, if we observed the simulated sequence $\{x_i\}_{i=1}^7$

0.7308 0.6513 0.5940 0.8049 0.7793 0.3342 0.2604

then we would stop on the fourth observation, since $x_4 = 0.8049 \ge v_3 = 0.6953$ is the first time our stopping condition is met. Hence, our gain is 0.8049, the value we stopped on. This is only slightly less than the expected gain of 0.8203 we saw above.

Example 3.2. This time, let $X_1, \dots, X_n$ be iid exponential random variables with mean $1/\lambda$.

These random variables have common pdf $f(x) = \lambda e^{-\lambda x}$ if $x \geq 0$ and $= 0$ elsewhere. Then, like above,

$$\begin{aligned} v_i &= E(\max\{X_{n-i+1}, v_{i-1}\}) \\ &= \int_0^\infty \lambda e^{-\lambda x} \max\{x, v_{i-1}\}\, dx \\ &= \int_0^{v_{i-1}} \lambda v_{i-1} e^{-\lambda x}\, dx + \int_{v_{i-1}}^\infty \lambda x e^{-\lambda x}\, dx \\ &= v_{i-1} + \frac{1}{\lambda} e^{-\lambda v_{i-1}} \end{aligned}$$

for $i = 1, \dots, n$.

Similarly to the previous example, letting $\lambda = 1$ (so $v_i = v_{i-1} + e^{-v_{i-1}}$) and $n = 7$ gives the values

n	0	1	2	3	4	5	6	7
v_i	0.0000	1.0000	1.3679	1.6225	1.8199	1.9820	2.1198	2.2398

with the expected reward being $v_7 = 2.2398$.

We have the same optimal stopping rule τ^* as in Example 3.1, and so for some simulated sequence $\{x_i\}_{i=1}^7$

$$0.1719 \quad 0.8306 \quad 2.0474 \quad 0.5930 \quad 3.4561 \quad 0.2426 \quad 1.1378$$

then we would stop on the third observation, since $x_3 = 2.0474 \geq v_4 = 1.8199$ is the first time our stopping condition is met.

3.2 Order Statistics

Order statistics are sample values placed in ascending order. In other words, the kth order statistic is equal to the kth-smallest value. For the purposes of this report, we are concerned only with the largest order statistic (called the nth order statistic), as we wish to compare our stopping rule with the maximum of the sequence.

In cases with a defined sample size of n, the nth order statistic, denoted as $X_{(n)}$, is the maximum of all the random variables in the sequence $X_1, \dots, X_n$. That is, $X_{(n)} = \max\{X_1, \dots, X_n\}$. The cumulative distribution function (cdf) of this statistic, assuming at least that each of the $\{X_i\}$ are independent, is

$$\begin{aligned} F_{X_{(n)}}(x) &= \mathrm{P}(\max\{X_1, \dots, X_n\} \leq x) \\ &= \mathrm{P}(X_1 \leq x, \dots, X_n \leq x) \\ &= \prod_{i=1}^n \mathrm{P}(X_i \leq x) \text{ by independence.} \end{aligned}$$

Taking this further for iid random variables, we obtain the nice expression

$$F_{X_{(n)}}(x) = \{\mathrm{P}(X \leq x)\}^n = \{F(x)\}^n,$$

where $F(x)$ is the common cdf of each X_i. The pdf of the nth order statistic is then $f_{X_{(n)}}(x) = nf(x)\{F(x)\}^{n-1}$, where $f(x)$ is the common pdf of each X_i. This follows from the relation between cdf and pdf, $f(x) = \frac{d}{dx}F(x)$.

4 Prophet Inequalities

Consider a gambler playing a game, where they must attempt to choose the highest value in a sequence of cards and they receive a reward equal to that value. They can use only non-anticipating strategies (i.e. rely only on the present and past) in order to choose the best time to stop on a value (higher values are better). Assuming the game is similar to the earlier idea of observing values sequentially with an inability to return to previous values, then we may apply our stopping rule in order to find the gambler's expected reward. We will define this reward as V, so that

$$V = \sup_{t \in T_n} E(X_t).$$

On the other hand, there may be another player of this game who always manages to win; they may be thought of as a cheater. More formally, we may call this person a "prophet", where they possess the ability of complete foresight, and so they may look ahead in the sequence and simply select the maximum value to always win. We can use the nth order statistic to represent this player, and hence define the prophet's expected reward, M, to be

$$M = E(\max\{X_1, \dots, X_n\}) = E(X_{(n)}).$$

A prophet inequality is then an inequality in M and V that holds for some class of sequences $X_1, \dots, X_n$. The immediate inequality is $M \geq V$; the gambler can never outplay the prophet. This inequality is quite obvious, so instead we may be curious as to how much better the prophet will perform compared to the gambler; this may be in both a multiplicative and additive sense.

4.1 Ratio Inequalities

It is perhaps intuitive that there should exist some $k \in \mathbb{R}$ so that k times the gambler's expected reward will outweigh the prophet's expected reward (i.e. $M \leq kV$) in our game above; otherwise the prophet's expected reward would be infinitely bigger than the gambler's expected reward. This value of k has enjoyed improvement over time, with initial findings bounding $k \in [2, 4]$ before the following result was derived [6].

Theorem 4.1. *If $X_1, ..., X_n$ are independent non-negative random variables, then*

$$M \leq 2V,$$

where the bound is sharp.

Proof. There exist many different proofs of this theorem. In this paper, we detail a threshold approach [8]. Define a fixed parameter, t, as a threshold such that we accept the first value $x_i \geq t$. To show the inequality, we will choose t to be the median of the distribution of the maximum (i.e. $P(\max_i X_i \geq t) = 1/2$) and define $x = P(\text{reject all rewards}) = P(\max_i X_i < t)$. Note that the strategy we are choosing is sub-optimal since there is the possibility of accepting no reward, however we can prove the theorem using it.

We first derive an upper bound on M. Letting $x^+ = \max\{0, x\}$,

$$M = E(\max_i X_i) \leq t + E[\max_i (X_i - t)^+] \leq t + \sum_i E[(X_i - t)^+].$$

Next a lower bound for the gambler's payoff (note that this is not V since we are using a sub-optimal rule; V has expected payoff at least as good as this), denoted as $E(\text{Prize})$:

$$V \geq E(\text{Prize}) \geq (1-x)t + x\sum_i E[(X_i - t)^+].$$

Now, it is clear that $x = P(\max_i X_i < t) = P(\max_i X_i \geq t) = 1/2$ since t is the median. Substituting this shows that

$$\begin{aligned} E(\text{Prize}) &\geq \frac{1}{2}\left(t + \sum_i E[(X_i - t)^+]\right) \\ &\geq \frac{1}{2}M, \end{aligned}$$

and hence we arrive at the desired result $M \leq 2V$.

To show that this bound is sharp, consider an example. Define the random variables $X_1 \equiv 1$ and $X_2 = 1/a$ with probability a for $a \in (0, 1]$, and 0 otherwise, and $X_n \equiv 0$ for all $n > 2$.
To find M, we need to consider when $X_1 \geq X_2$ and when $X_1 < X_2$. So, for $X_1 \geq X_2$, we require $X_2 = 0$, which occurs with probability $1 - a$, and hence $X_1 < X_2$ occurs with probability a. Thus

$$\begin{aligned} M &= E(\max_i X_i) \\ &= P(X_1 \geq X_2) \cdot E(X_1 \mid X_1 \geq X_2) + P(X_1 < X_2) \cdot E(X_2 \mid X_1 < X_2) \\ &= (1-a) \cdot 1 + a \cdot \frac{1}{a} \\ &= 2 - a. \end{aligned}$$

Now, for V we simply have to determine which random variable has a higher expected value. Clearly $E(X_1) = 1$, but we also have $E(X_2) = a \cdot \frac{1}{a} = 1$, so

$$V = \sup_{t \in T_n} E(X_t) = 1.$$

Then as $a \to 0$, the limiting equality is apparent, and hence the bound is sharp.
Interestingly, the bound is only sharp in particular cases, such as in this "long-shot" example; more information can be found in [4]. □

Theorem 4.1 tells us that, on average, the prophet cannot win more than twice that of the gambler, which may be quite surprising. Both bounded and unbounded sequences of random variables satisfying the above assumptions abide by the inequality, such as the uniform and exponential distribution examples we saw earlier. If we were to relax the assumptions, we arrive at the following results (both coming from [4]):

- If the independence assumption is dropped, then $M \leq nV$, where n is the number of rvs in the sequence;

- If the non-negativity assumption is dropped, then the ratio may be arbitrary.

We can extend Theorem 4.1 to the possibility of obtaining negative values whilst maintaining a meaningful result by defining $X^+ = \max\{0, X\}$, replacing any negative values with 0; this would allow something like stock prices or offers on a house to be modelled using the theorem.

4.2 Difference Inequalities

We now wish to consider results of the form $M - V \leq p$, allowing us to compare the additive advantage the prophet enjoys over the gambler. Before stating and proving the appropriate theorem, we shall define some notations and introduce lemmas that will aid our proof, with all original work from [3].

Definition 4.2. For random variables $X_1, \dots, X_n$, define $V(X_j, \dots, X_n) = E(\max\{X_j, V(X_{j+1}, \dots, X_n)\})$ for $j = 1, \dots, n-1$ (this is analogous to $v_j = E(\max\{X_j, v_{j-1}\})$), and define

$$D(X_1, \dots, X_n) = E(\max\{X_1, \dots, X_n\}) - V(X_1, \dots, X_n),$$

the additive advantage of the prophet over the gambler.

Definition 4.3. For an integrable random variable Y with pdf $f(y)$ and constants $-\infty < a < b < \infty$, let Y_a^b denote a random variable with

$$Y_a^b = \begin{cases} Y, & \text{with probability } 1 - \int_a^b f(y)\,dy, \\ a, & \text{with probability } (b-a)^{-1} \int_a^b (b-y) f(y)\,dy, \\ b, & \text{with probability } (b-a)^{-1} \int_a^b (y-a) f(y)\,dy. \end{cases}$$

and say that Y_a^b is extremal with respect to Y, a, and b.

Lemma 4.4. *Using Y_a^b as defined in Definition 4.3, we have that $E(Y_a^b) = E(Y)$.*

A proof is supplied in the appendix.

Lemma 4.5. *Let Y, a, b, and Y_a^b be as above. If X is any integrable random variable independent of both Y and Y_a^b, then $E(\max\{X, Y\}) \leq E(\max\{X, Y_a^b\})$.*

A proof of this lemma is supplied in the appendix.

Lemma 4.6. *Given $n > 2$ and independent random variables $X_1, \dots, X_n$, there exists a zero-one valued random variable W independent of $X_2, \dots, X_{n-2}$, satisfying $D(X_1, \dots, X_n) \leq D(\mu, X_2, \dots, X_{n-2}, W)$, where $\mu = V(X_2, \dots, X_n)$.*

Proof. Because this proof is central to the proof of the next theorem, we include the proof here instead of the appendix.
Without loss of generality, assume that all random variables are bounded in [0,1]. Adding or multiplying by

suitable constants is allowed to achieve a general result. By Definition 4.2, and noting that $V(\mu, X_2, \dots, X_n) = V(X_2, \dots, X_n) = \mu$,

$$\begin{aligned} D(X_1, \dots, X_n) &= E(\max\{X_1, \dots, X_n\}) - V(X_1, \dots, X_n) \\ &\leq E(\max\{\mu, X_2, \dots, X_n\}) + E(X_1 - \mu)^+ - E(\max\{X_1, V(X_2, \dots, X_n)\}) \\ &= E(\max\{\mu, X_2, \dots, X_n\}) + E(X_1 - \mu)^+ - V(X_2, \dots, X_n) - E(X_1 - V(X_2, \dots, X_n)^+ \\ &= D(\mu, X_2, \dots, X_n). \end{aligned} \tag{1}$$

Now define the extremal random variables $Z = (X_n)_0^1$ and $Y = (X_{n-1})_{E(X_n)}^1$, letting them be independent of each other and of $X_1, \dots, X_{n-2}$. Using Lemma 4.4, we know that $E(Z) = E(X_n)$ and $E(Y) = E(X_{n-1})$. Further, we see that

$$\begin{aligned} V(X_{n-1}, X_n) &= E(\max\{X_{n-1}, E(X_n)\}) \\ &= E(X_n) \cdot P(X_{n-1} \leq E(X_n)) + E(X_{n-1}) \cdot P(X_{n-1} > E(X_n)) \end{aligned}$$

and

$$\begin{aligned} V(Y, Z) &= E(\max\{Y, E(Z)\}) \\ &= E(Z) \cdot P(Y \leq E(Z)) + E(Y) \cdot P(Y > E(Z)) \\ &= E(X_n) \cdot P(Y \leq E(X_n)) + E(X_{n-1}) \cdot P(Y > E(X_n)) \\ &= E(X_n) \cdot P(X_{n-1} \leq E(X_n)) + E(X_{n-1}) \cdot P(X_{n-1} > E(X_n)) \\ &= V(X_{n-1}, X_n), \end{aligned}$$

and thus $V(\mu, X_2, \dots, X_n) = V(\mu, X_2, \dots, X_{n-2}, Y, Z)$.
By Lemma 4.5, $E(\max\{\mu, X_2, \dots, X_n\}) \leq E(\max\{\mu, X_2, \dots, X_{n-2}, Y, Z\})$, and thus

$$D(\mu, X_2, \dots, X_n) \leq D(\mu, X_2, \dots, X_{n-2}, Y, Z). \tag{2}$$

Finally, letting W be a random variable independent of $X_2, \dots, X_{n-2}$ where $W = 1$ with probability $V(Y, Z)$ and $= 0$ otherwise, we have that

$$D(\mu, X_2, \dots, X_{n-2}, Y, Z) = D(\mu, X_2, \dots, X_{n-2}, W), \tag{3}$$

because $E(W) = V(Y, Z)$ and $E(Z) = E(X_n) \leq \mu$.
Combining (1), (2), and (3) gives the required result. $\square$

We may now introduce the main result and its proof.

Theorem 4.7. *If $X_1, \dots, X_n$ are independent random variables bounded in [a,b], then $M - V \leq (b-a)/4$, and this bound is best possible.*

Proof. Again, without loss of generality, assume that all random variables are bounded in $[0, 1]$.
Starting with a sequence of independent random variables $X_1, \ldots, X_n$, we may apply Lemma 4.6 recursively to reduce to the case of considering only X_1 and X_2.
Using the same definition for μ as above, let $E(X_2) = \mu$ and introduce a random variable $Z = 1$ with probability μ and $= 0$ otherwise. Then

$$\begin{aligned} D(X_1, X_2) &\leq D(\mu, Z) \\ &= E(\max\{\mu, Z\}) - V(\mu, Z) \\ &= \{\mu \cdot P(Z \leq \mu) + E(Z) \cdot P(Z > \mu)\} - E(\max\{\mu, E(Z)\}) \\ &= (\mu(1-\mu) + \mu) - \mu \\ &= \mu - \mu^2 \\ &\leq 1/4, \end{aligned}$$

since $\mu \in [0, 1]$. The bound is attained when $\mu = 1/2$, so the inequality is sharp. □

4.3 The Prophet Region

This last result is interesting and allows the above results to be simple corollaries for the case of each sequence of $X_1, \ldots, X_n$ being independent and bounded in $[0, 1]$. The result is a two-sided inequality bounding the prophet's possible outcomes relative to functions of the gambler's possible values.

Definition 4.8. A convex set is such that for any two points in the set, their line segment is contained within the set. Intuitively, you may imagine a circle or cube to be a convex set, whereas shapes with indents (e.g. a crescent) are not convex. A closed convex set is then as above, but with the addition that all the limit points are contained within the set.

Theorem 4.9. *The set of all possible values of the ordered pair (V, M) for independent random variables $X_1, \ldots, X_n$ bounded in $[0, 1]$ is precisely the closed convex set S in $\mathbb{R}^2$ given by*

$$S = \left\{(x, y) \ : \ x \leq y \leq 2x - x^2, \ 0 \leq x \leq 1\right\}.$$

Proof. As the theorem is a precise result, we must prove that both the statement "Every point in S is a possible value of (V, M)" and its converse hold. First consider the original statement.
We can show the first statement by example. Start by letting $\alpha = \frac{y-x}{x-x^2}$. Define $X_1, \ldots, X_n$ by $X_1 \equiv x, X_2 = 1$ with probability αx and $= 0$ otherwise, and $X_n = 0$ for all $n > 2$. We can now check that in fact $x = V$ and $y = M$.
Note that $E(X_1) = x, \ E(X_2) = \alpha x$ and $E(X_n) = 0$. Also note that $0 \leq \alpha \leq 1 \Rightarrow \alpha x \leq x$, since $y \leq 2x - x^2$.
Then by definition, $V = \sup_{t \in T_n}\{E(X_t)\} = x$.
Next we need to find when $X_1 > X_2$ and when $X_2 > X_1$ to determine the distribution of $X_{(n)}$. We have that

$X_1 > X_2$ when $X_2 = 0$, which occurs with probability $1 - \alpha x$, and $X_2 > X_1$ when $X_2 = 1$, which occurs with probability αx. Hence $X_{(n)} = 1$ with probability αx and $= x$ otherwise. Then

$$\begin{aligned} M = E(X_{(n)}) &= x(1-\alpha x) + \alpha x \\ &= \alpha x + x - \alpha x^2 \\ &= \alpha(x - x^2) + x \\ &= y - x + x \\ &= y. \end{aligned}$$

Hence every point in S is a possible value of (V, M).

Now to prove the converse; every possible value of (V, M) is contained in S.
The proof here will run very similarly to the proof of Theorem 4.7. Using Lemma 4.6, we may reduce $X_1, \ldots, X_n$ down to considering only X_1, X_2 by defining the random variables $Y_1 \equiv V$ and $Y_2 = 1$ with probability V, and $= 0$ otherwise. Before continuing, we note that

$$\begin{aligned} E(\max\{Y_1, Y_2\}) &= P(Y_1 \geq Y_2) \cdot E(Y_1 \mid Y_1 \geq Y_2) + P(Y_1 < Y_2) \cdot E(Y_2 \mid Y_1 < Y_2) \\ &= (1 - V) \cdot V + V \cdot 1 \\ &= 2V - V^2 \end{aligned}$$

and $\sup_{t \in T_2} E(Y_t) = V$.
Then applying Lemma 4.6,

$$\begin{aligned} D(X_1, X_2) = M - V \leq D(Y_1, Y_2)& \\ &= E(\max\{Y_1, Y_2\}) - \sup_{t \in T_2} E(Y_t) \\ &= V - V^2. \end{aligned}$$

Hence we have the inequality $M \leq 2V - V^2$. The inequality $M \geq V$ is apparent from the definitions of M and V, and so we arrive at the desired result $V \leq M \leq 2V - V^2$. □

5 Simulations and Results

The following examples demonstrate the above results empirically for a variety of situations. The standard uniform distribution is bounded on $[0, 1]$ so all the above results apply, whereas the exponential distribution is unbounded, and so only the ratio result will necessarily hold (although we will derive an interesting difference result through these simulations). We begin by deriving the theoretical results for the prophet's and gambler's expected rewards before running simulations to compare results.

5.1 The $U(0,1)$ Distribution

Let $X_1, \dots, X_n$ be a sequence of iid $U(0,1)$ random variables. Then their common pdf is $f(x) = 1$ on $[0,1]$ and cdf is $F(x) = x$ on $[0,1]$. Hence the cdf of the nth order statistic is $F_{X_{(n)}}(x) = x^n$ for $x \in [0,1]$. This has pdf $f_{X_{(n)}}(x) = nx^{n-1}$ and we can then calculate the expected value

$$\begin{aligned} M = E(X_{(n)}) &= \int_{-\infty}^{\infty} x \cdot f_{X_{(n)}}(x)\, dx \\ &= \int_0^1 x \cdot nx^{n-1}\, dx \\ &= \frac{n}{n+1}. \end{aligned}$$

We can then derive the expected reward for each n for the gambler, which was done in Example 3.1. Generating the table and plots gives us:

n	1	2	3	$\cdots$	48	49	50
M	0.5000	0.6667	0.7500	$\cdots$	0.9796	0.9800	0.9804
V	0.5000	0.6250	0.6953	$\cdots$	0.9628	0.9635	0.9641

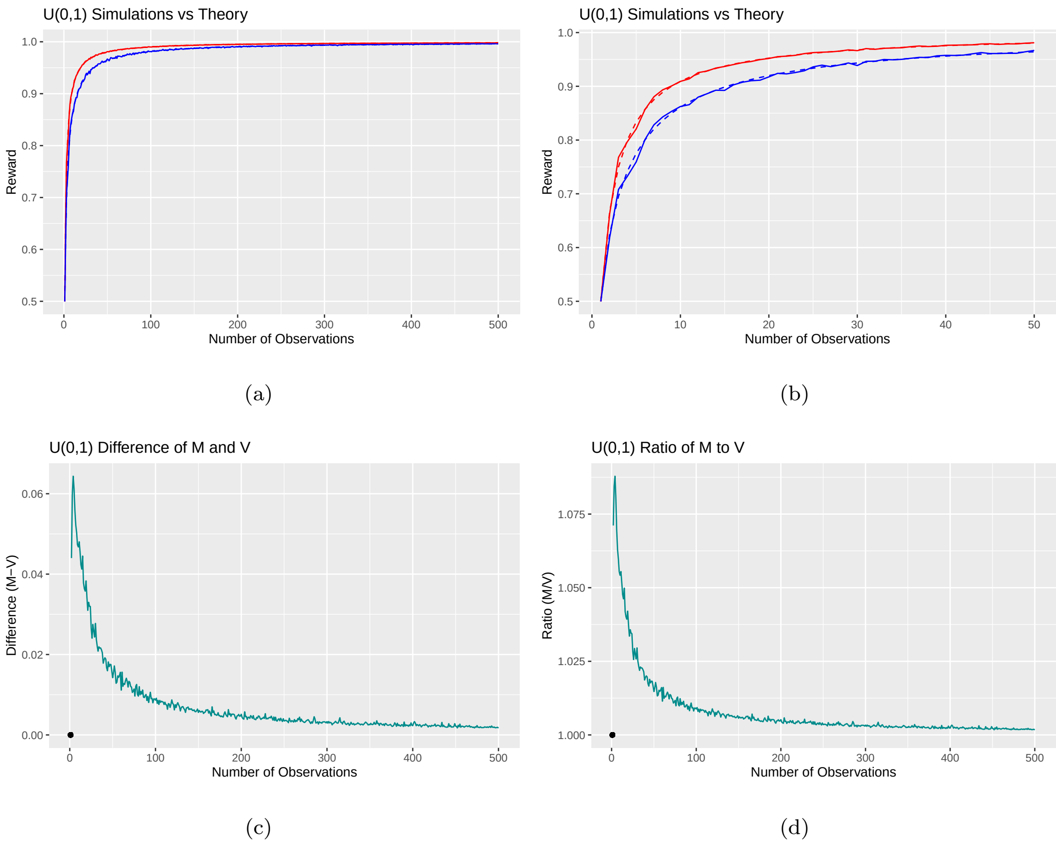

(a) (b) (c) (d)

Figure 1: Plots demonstrating each prophet inequality for the $U(0,1)$ distribution. In (a) and (b), the red line denotes the prophet and the blue line the gambler, whilst the solid lines are the simulated averages and the dashed lines the theoretical results.

Plots (a) and (b) in Figure 1 show how the simulations agree with the theoretical results, both in the sense of the individual prophet and gambler results following their theory, but also that the prophet's reward is always at least as large as the gamblers (i.e. $M \geq V$). Plots (c) and (d) demonstrate the difference, $M - V$, and ratio, M/V, respectively, and comparing the results to our inequalities, we see they fall comfortably below the bounds; in many cases (such as this), the bounds will be quite conservative. Additionally, we can see a clear inverse relationship between them and n. This would indicate a possible asymptotic result, and we can show that this is in fact the case. It can be shown [2] that V is asymptotically $b - 2(b-a)/n$, which approaches b as $n \to \infty$. Hence $V \to b = 1$ in our example. On the other hand, using our expression for M it is clear that $M \to 1$ as $n \to \infty$ also, and hence $M - V \to 0$ and $M/V \to 1$. Both of these observations are consistent with the plots.

5.2 The $Exp(1)$ Distribution

Let $X_1, \ldots, X_n$ be a sequence of iid $Exp(1)$ random variables. They have common pdf $f(x) = e^{-x}$ and cdf $F(x) = 1 - e^{-x}$ for $x \geq 0$. Hence the cdf of the nth order statistic is $F_{X_{(n)}}(x) = (1 - e^{-x})^n$ and so $f_{X_{(n)}}(x) = ne^{-x}(1 - e^{-x})^{n-1}$ for $x \geq 0$.

Now, we will not use the traditional definition for expectation in this example as the integral is quite difficult to solve. Instead, we turn to a result that will allow us to calculate the expected value recursively.

Lemma 5.1. *Let X be a non-negative continuous integrable random variable. Then its expectation can be evaluated using*

$$E(X) = \int_0^\infty (1 - F(x))\, dx.$$

The proof is provided in the appendix.

Let us define $E_n = E(X_{(n)})$ so that using Lemma 5.1, we have

$$\begin{aligned} E_n - E_{n-1} &= \int_0^\infty \left(1 - (1 - e^{-x})^n\right)\, dx - \int_0^\infty \left(1 - (1 - e^{-x})^{n-1}\right)\, dx \\ &= \int_0^\infty (1 - e^{-x})^{n-1}[1 - (1 - e^{-x})]\, dx \\ &= \int_0^\infty (1 - e^{-x})^{n-1} e^{-x}\, dx \\ &= 1/n. \end{aligned}$$

Then continuing in a similar fashion,

$$\begin{aligned} E_n &= (E_n - E_{n-1}) + (E_{n-1} - E_{n-2}) + \cdots + (E_2 - E_1) + E_1 \\ &= \frac{1}{n} + \frac{1}{n-1} + \cdots + \frac{1}{2} + 1 \\ &= H_n, \end{aligned}$$

where H_n denotes the nth partial sum of the harmonic series.

We can now produce a similar table encapsulating the values for the gambler and prophet for different n values, along with the plots showing the raw averages and the difference and ratios, like for the $U(0, 1)$ example.

n	1	2	3	$\cdots$	48	49	50
M	1.0000	1.5000	1.8333	$\cdots$	4.4588	4.4792	4.4992
V	1.0000	1.3679	1.6225	$\cdots$	3.9368	3.9563	3.9755

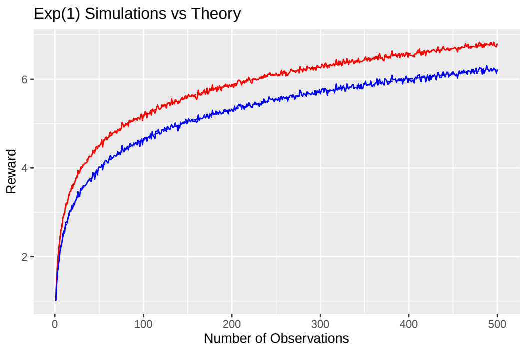

(a)

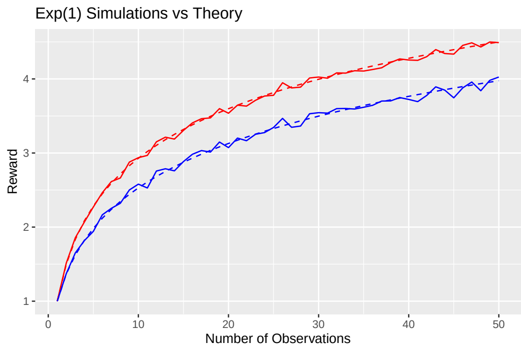

(b)

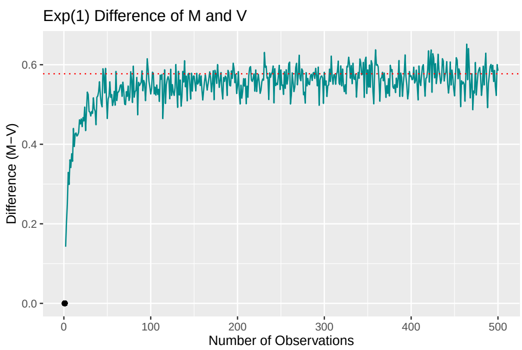

(c)

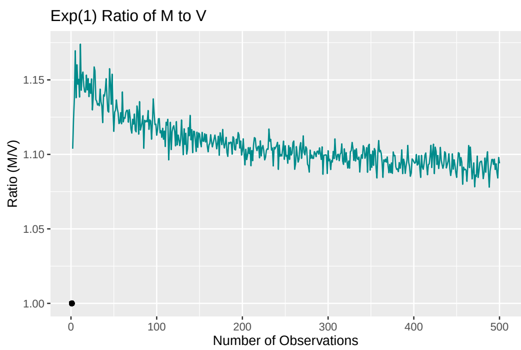

(d)

Figure 2: Plots demonstrating each prophet inequality for the $Exp(1)$ distribution. In (a) and (b), the red line denotes the prophet and the blue line the gambler, whilst the solid lines are the simulated averages and the dashed lines the theoretical results.

Plots (a) and (b) in Figure 2 demonstrate how the experimental results closely follow the theory as expected, with deviations occurring only due to the nature of probability. Looking at (d), we see again that the ratio, M/V, lies well below 2, indicating another conservative result. On the other hand, our difference result (Theorem 4.7) may not be applied since our random variables are unbounded. Hence we currently do not know whether the difference, $M - V$, is bounded or not. Looking at (c), we see that the difference does seem to converge around the red dotted line, however to confirm this we will analyse the asymptotic results for M and V. It will be shown in Appendix B that, in fact, we have the asymptotic result $M - V \to \gamma \approx 0.5772$, where γ is the Euler-Mascheroni constant.

6 Future Work

Optimal stopping theory contains a broad spectrum of possible future research for prophet inequalities. We can extend the simple case examined in this report of stopping only once into cases of multiple stopping, where you can stop and accept multiple rewards along the sequence. This extends the reward function; letting $v^{n,k}$ denote the value of a game with k stops and n observations left to make, then

$$\begin{aligned}
v^{n,1} &= E(\max\{X_{N-n+1}, v^{n-1,1}\}), \quad 1 \le n \le N, \\
v^{0,1} &= -\infty, \\
v^{n,k-i+1} &= E(\max\{X_{N-n+1} + v^{n-1,k-i}, v^{n-1,k-i+1}\}), \quad k-i+1 \le n \le N, \\
v^{k-i,k-i+1} &= -\infty, \quad i = k-1, \dots, 1,
\end{aligned}$$

and we naturally obtain a multi-dimensional stopping rule as well.

Otherwise, many of the assumptions made in the theorems can be varied to achieve different results. For example, we may relax the independence and non-negativity assumptions, assume the distributions are known/unknown, or further tighten assumptions to assume iid in order to develop new results. All of these have been studied previously, however the more complex case of multiple stopping detailed above has not received much attention.

7 Conclusion

This report successfully updates terminology and expands on prior papers on prophet inequalities in order to improve readability and initial comprehensions of the material. We were able to understand where the inequalities come from, what they mean from a physical perspective, and their place in the world with room for further research in situations of multiple stopping and changing initial assumptions. Demonstrating the results empirically allows a visualisation of the theorems to see how actual results develop over time and having different numbers of values available to observe.

References

[1] Morris H DeGroot. *Optimal statistical decisions.* John Wiley & Sons, 2005.

[2] Hugh N Entwistle, Christopher J Lustri, and Georgy Yu Sofronov. On asymptotics of optimal stopping times. *Mathematics*, 10(2):194, 2022.

[3] Theodore P Hill and Robert P Kertz. Additive comparisons of stop rule and supremum expectations of uniformly bounded independent random variables. *Proceedings of the American Mathematical Society*, 83(3):582–585, 1981.

[4] Theodore P Hill and Robert P Kertz. Ratio comparisons of supremum and stop rule expectations. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 56(2):283–285, 1981.

[5] Theodore P Hill and Robert P Kertz. A survey of prophet inequalities in optimal stopping theory. *Contemp. Math*, 125:191–207, 1992.

[6] Krengel, Sucheston, and Garling. On semiamarts, amarts, and processes with finite value. *Advances in Applied Probability*, 4:197–266, 1978.

[7] Georgy Sofronov. An optimal sequential procedure for a multiple selling problem with independent observations. *European Journal of Operational Research*, 225(2):332–336, 2013.

[8] Subhash Suri. Optimal stopping rules. `https://sites.cs.ucsb.edu/~suri/ccs130a/OptStopping.pdf`, 2020.

A Simulation Code

U(0,1) code:

```
# U(0,1) Simulations and Theoretical Results
library(stats)
library(ggplot2)

N = seq(1, 500) # number of observations of random variables

# Theory
v = zeros(length(N)+1, 1)
m = zeros(length(N)+1, 1)
v[1] = 0 # expected reward with no observations left since R indexes from 1
m[1] = 0 # prophet reward with no observations left

for(i in 1:length(N)){
  v[i+1] = (1+v[i]^2)/2 # gambler expected reward for each number of obs left
  m[i+1] = N[i]/(N[i]+1) # prophet expected reward for each number of obs left
}
# remember m[i] means there are i+1 observations left

# Simulations
n = 1000 # number of simulations
V = zeros(n, length(N)) # matrix to store result of each sim for each num of rv
M = zeros(n, length(N))

V_raw = zeros(n, 1) # results from sim
M_raw = zeros(n, 1)

V_avg = zeros(length(N), 1) # average of sim for each num of obs
M_avg = zeros(length(N), 1)

V_hold = 0 # initialising variable

for(p in 1:length(N)){
  for(q in 1:n){
    x = runif(p, 0, 1)
    for(r in 1:p){
      if(x[r] >= v[p-r+1]){
        V_hold = x[r]
        break
      }
    }
    V_raw[q] = V_hold
    M_raw[q] = max(x)
  }
```

```
  V_avg[p] = mean(V_raw)
  M_avg[p] = mean(M_raw)
}

V_thy = v[2:501]
M_thy = m[2:501]

#----------------------

avg_df = data.frame(N, M_avg, V_avg)

sim_vs_thy_plot = ggplot(data = avg_df, aes(x = N)) +
  geom_line(aes(y = M_avg, colour = "Prophet")) +
  geom_line(aes(y = V_avg, colour = "Gambler")) +
  geom_line(aes(y = M_thy, colour = "Prophet"), linetype = "dashed") +
  geom_line(aes(y = V_thy, colour = "Gambler"), linetype = "dashed") +
  scale_colour_manual("", breaks = c("Prophet", "Gambler"), values = c("red", "blue")) +
  xlab("Number of Observations") +
  scale_y_continuous("Reward") +
  labs(title = "U(0,1) Simulations vs Theory") +
  theme(legend.position = "none")

avg_ratio = M_avg/V_avg
ratio_df = data.frame(seq(1, 500), avg_ratio)
ratio_df[1,2] = NA

ratio_plot = ggplot(data = ratio_df, aes(x = N)) +
  geom_line(aes(y = avg_ratio), colour = "darkcyan") +
  xlab("Number of Observations") +
  scale_y_continuous("Ratio (M/V)") +
  labs(title = "U(0,1) Ratio of M to V") +
  geom_point(aes(x = 1, y=1))

avg_diff = M_avg - V_avg
diff_df = data.frame(seq(1, 500), avg_diff)
diff_df[1,2] = NA

diff_plot = ggplot(data = diff_df, aes(x = N)) +
  geom_line(aes(y = avg_diff), colour = "darkcyan") +
  xlab("Number of Observations") +
  scale_y_continuous("Difference (M-V)") +
  labs(title = "U(0,1) Difference of M and V") +
  geom_point(aes(x = 1, y = 0))

# smaller set (to show differences more easily)
avg_sml = avg_df[1:50,]
```

```
M_thy_sml = M_thy[1:50]
V_thy_sml = V_thy[1:50]

sim_plot_sml = sim_vs_thy_plot = ggplot(data = avg_sml, aes(x = seq(1,50))) +
  geom_line(aes(y = M_avg[1:50], colour = "Prophet",)) +
  geom_line(aes(y = V_avg[1:50], colour = "Gambler",)) +
  geom_line(aes(y = M_thy_sml, colour = "Prophet"), linetype = "dashed") +
  geom_line(aes(y = V_thy_sml, colour = "Gambler"), linetype = "dashed") +
  scale_colour_manual("", breaks = c("Prophet", "Gambler"), values = c("red", "blue")) +
  xlab("Number of Observations") +
  scale_y_continuous("Reward") +
  labs(title = "U(0,1) Simulations vs Theory") +
  theme(legend.position = "none")

avg_extra_sml = avg_df[1:20,]
M_thy_extra_sml = M_thy[1:20]
V_thy_extra_sml = V_thy[1:20]

sim_plot_extra_sml = ggplot(data = avg_extra_sml, aes(x = seq(1,20))) +
  geom_line(aes(y = M_avg[1:20], colour = "Prophet",)) +
  geom_line(aes(y = V_avg[1:20], colour = "Gambler",)) +
  geom_line(aes(y = M_thy_extra_sml, colour = "Prophet"), linetype = "dashed") +
  geom_line(aes(y = V_thy_extra_sml, colour = "Gambler"), linetype = "dashed") +
  scale_colour_manual("", breaks = c("Prophet", "Gambler"), values = c("red", "blue")) +
  xlab("Number of Observations") +
  scale_y_continuous("Reward") +
  labs(title = "U(0,1) Simulations vs Theory") +
  theme(legend.position = "none")
```

Exp(1) code:

```
# Exp(1) Simulations and Theoretical Results
library(stats)
library(ggplot2)

N = seq(1, 500) # number of observations of random variables

# Theory
v = zeros(length(N)+1, 1)
m = zeros(length(N)+1, 1)
v[1] = 0 # expected reward with no observations left since R indexes from 1
m[1] = 0 # prophet reward with no observations left

for(i in 1:length(N)){
  v[i+1] = v[i] + exp(-v[i]) # gambler expected reward for each number of obs left
  m[i+1] =  1/N[i] + m[i]  # prophet expected reward for each number of obs left
}
# remember m[i] means there are i+1 observations left

# Simulations
n = 1000 # number of simulations
V = zeros(n, length(N)) # matrix to store result of each sim for each num of rv
M = zeros(n, length(N))

V_raw = zeros(n, 1) # results from sim
M_raw = zeros(n, 1)

V_avg = zeros(length(N), 1) # average of sim for each num of obs
M_avg = zeros(length(N), 1)

V_hold = 0 # initialising variable

for(p in 1:length(N)){
  for(q in 1:n){
    x = rexp(p, 1)
    for(r in 1:p){
      if(x[r] >= v[p-r+1]){
        V_hold = x[r]
        break
      }
    }
    V_raw[q] = V_hold
    M_raw[q] = max(x)
  }
  V_avg[p] = mean(V_raw)
  M_avg[p] = mean(M_raw)
```

```
}

V_thy = v[2:501]
M_thy = m[2:501]

#-----------------------

avg_df = data.frame(N, M_avg, V_avg)

sim_vs_thy_plot = ggplot(data = avg_df, aes(x = N)) +
  geom_line(aes(y = M_avg, colour = "Prophet")) +
  geom_line(aes(y = V_avg, colour = "Gambler")) +
  geom_line(aes(y = M_thy, colour = "Prophet"), linetype = "dashed") +
  geom_line(aes(y = V_thy, colour = "Gambler"), linetype = "dashed") +
  scale_colour_manual("", breaks = c("Prophet", "Gambler"), values = c("red", "blue")) +
  xlab("Number of Observations") +
  scale_y_continuous("Reward") +
  labs(title = "Exp(1) Simulations vs Theory") +
  theme(legend.position = "none")

avg_ratio = M_avg/V_avg
ratio_df = data.frame(seq(1, 500), avg_ratio)
ratio_df[1,2] = NA

ratio_plot = ggplot(data = ratio_df, aes(x = N)) +
  geom_line(aes(y = avg_ratio), colour = "darkcyan") +
  xlab("Number of Observations") +
  scale_y_continuous("Ratio (M/V)") +
  labs(title = "Exp(1) Ratio of M to V") +
  geom_point(aes(x = 1, y=1))

avg_diff = M_avg - V_avg
diff_df = data.frame(seq(1, 500), avg_diff)
diff_df[1,2] = NA

diff_plot = ggplot(data = diff_df, aes(x = N)) +
  geom_line(aes(y = avg_diff), colour = "darkcyan") +
  xlab("Number of Observations") +
  scale_y_continuous("Difference (M-V)") +
  labs(title = "Exp(1) Difference of M and V") +
  geom_point(aes(x = 1, y = 0)) +
  geom_hline(yintercept = 0.5772156649, linetype = "dotted", colour = "red")

# smaller set (to show differences more easily)
avg_sml = avg_df[1:50,]
M_thy_sml = M_thy[1:50]
```

```
V_thy_sml = V_thy[1:50]

sim_plot_sml = sim_vs_thy_plot = ggplot(data = avg_sml, aes(x = seq(1,50))) +
  geom_line(aes(y = M_avg[1:50], colour = "Prophet",)) +
  geom_line(aes(y = V_avg[1:50], colour = "Gambler",)) +
  geom_line(aes(y = M_thy_sml, colour = "Prophet"), linetype = "dashed") +
  geom_line(aes(y = V_thy_sml, colour = "Gambler"), linetype = "dashed") +
  scale_colour_manual("", breaks = c("Prophet", "Gambler"), values = c("red", "blue")) +
  xlab("Number of Observations") +
  scale_y_continuous("Reward") +
  labs(title = "Exp(1) Simulations vs Theory") +
  theme(legend.position = "none")

avg_extra_sml = avg_df[1:20,]
M_thy_extra_sml = M_thy[1:20]
V_thy_extra_sml = V_thy[1:20]

sim_plot_extra_sml = ggplot(data = avg_extra_sml, aes(x = seq(1,20))) +
  geom_line(aes(y = M_avg[1:20], colour = "Prophet",)) +
  geom_line(aes(y = V_avg[1:20], colour = "Gambler",)) +
  geom_line(aes(y = M_thy_extra_sml, colour = "Prophet"), linetype = "dashed") +
  geom_line(aes(y = V_thy_extra_sml, colour = "Gambler"), linetype = "dashed") +
  scale_colour_manual("", breaks = c("Prophet", "Gambler"), values = c("red", "blue")) +
  xlab("Number of Observations") +
  scale_y_continuous("Reward") +
  labs(title = "Exp(1) Simulations vs Theory") +
  theme(legend.position = "none")
```

B Proofs

B.1 Proof of $M - V$ Asymptotic for $Exp(1)$ Case

For this proof, let log denote the natural logarithm. It can be shown [2] that V is asymptotically $\beta \log(n)$, where $\beta = 1$ in our case. Further, we know from earlier that $M = H_n$. Now we wish to derive an asymptotic result for H_n.

First we note that

$$1 + \frac{1}{2} + \cdots + \frac{1}{n} > \int_1^n \frac{1}{x}\, dx = \log(n) \tag{4}$$

and

$$\frac{1}{2} + \cdots + \frac{1}{n} < \int_1^n \frac{1}{x}\, dx = \log(n). \tag{5}$$

These results may be justified geometrically via the following plot:

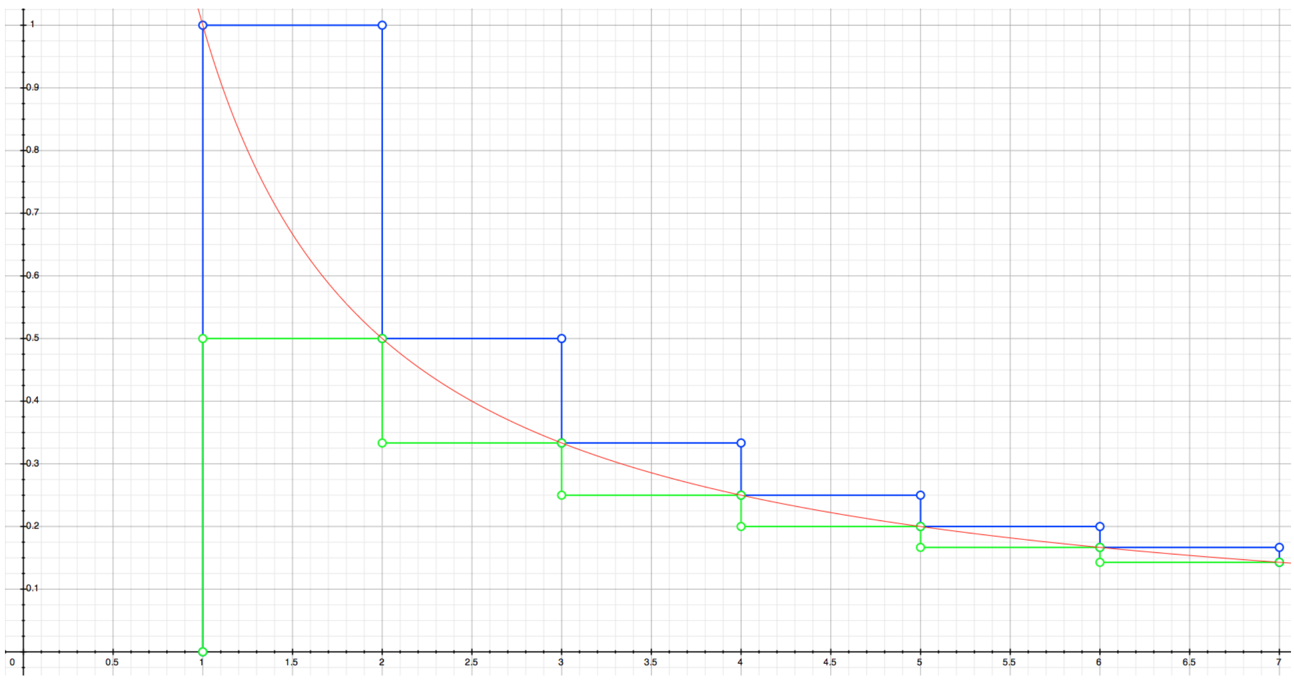

Figure 3: Geometric Argument for Bounds on H_n. [Source: `https://math.stackexchange.com/questions/156326/showing-inequality-for-harmonic-series`]

The area of the blue curve is an upper rectangle (left side) Riemann approximation representing the sum $\sum_{k=1}^{n} \frac{1}{k}$, whilst the green curve is a lower rectangle (right side) Riemann approximation representing the sum $\sum_{k=1}^{n} \frac{1}{k+1}$, which bounds the true area of the curve between these two values. Combining (4) and (5) gives

$$\log(n) < H_n < 1 + \log(n),\ \ n > 1.$$

What this shows is that the rate of growth of H_n is bounded by $\log(n)$. We know that $\log(n)$ diverges as $n \to \infty$ and so H_n also diverges.

Now, rearranging the inequality above gives

$$0 < H_n - \log(n) < 1.$$

If we let $a_n = H_n - \log(n)$, we can see that a_n is a monotonically increasing sequence and is bounded. Applying the monotonic convergence theorem (if a sequence increases/decreases and is bounded above/below by a supremum/infimum, it will converge), we conclude that $\lim_{n\to\infty} a_n$ exists. Computing this limit, we have

$$\lim_{n\to\infty} a_n = \gamma,$$

where $\gamma \approx 0.5772$ is the Euler-Mascheroni constant. Hence

$$H_n - \log(n) \approx \gamma$$

for large n.

As a result, we get that

$$\begin{aligned} M - V &\approx \gamma + \log(n) - \log(n) \\ &= \gamma \end{aligned}$$

for large n, as required. Going back to Figure 2, the red line in (c) is valued at γ, confirming this result.

B.2 Proof of Lemma 4.4:

First note that Y_a^b is a mixed random variable. Hence its cdf is the sum of continuous and step functions. Letting Y_a^b have cdf denoted by $F_Y(y)$, we have

$$F_Y(y) = C(y) + D(y),$$

where C denotes the continuous part and D the step part of the cdf. We can then differentiate C, denoted as c, and hence we have the property

$$\int_{-\infty}^{\infty} c(y)\,dy + \sum_{y_k} P(Y = y_k) = 1.$$

Thus we can obtain the expected value of Y_a^b as

$$E(Y_a^b) = \int_{-\infty}^{\infty} yc(y)\,dy + \sum_{y_k} y_k P(Y = y_k).$$

Applying this to our random variable, we obtain

$$\begin{aligned} E(Y_a^b) &= \int_{-\infty}^{a} yf(y)\,dy + \int_{b}^{\infty} yf(y)\,dy + aP(Y_a^b = a) + bP(Y_a^b = b) \\ &= \int_{-\infty}^{a} yf(y)\,dy + \int_{b}^{\infty} yf(y)\,dy + a(b-a)^{-1}\int_a^b (b-y)f(y)\,dy + b(b-a)^{-1}\int_a^b (y-a)f(y)\,dy \\ &= \int_{-\infty}^{a} yf(y)\,dy + \int_{b}^{\infty} yf(y)\,dy + (b-a)^{-1}\int_a^b (ab - ay + by - ab)f(y)\,dy \\ &= \int_{-\infty}^{\infty} yf(y)\,dy \\ &= E(Y). \end{aligned}$$

B.3 Proof of Lemma 4.5:

First define the random variable X that is independent of Y and Y_a^b. We are required to prove that

$$\int_a^b \max\{x,y\}f(y)\,dy \le (b-a)^{-1}\left\{\max\{x,a\}\int_a^b (b-y)f(y)\,dy + \max\{x,b\}\int_a^b (y-a)f(y)\,dy\right\}. \tag{6}$$

To see this, consider

- If $x \le a \le y \le b$, then

$$(b-a)^{-1}\left\{a\int_a^b (b-y)f(y)\,dy + b\int_a^b (y-a)f(y)\,dy\right\} = \int_a^b yf(y)\,dy.$$

- If $a \le x \le y \le b$, then

$$\begin{aligned}
&(b-a)^{-1}\left\{x\int_a^b (b-y)f(y)\,dy + b\int_a^b (y-a)f(y)\,dy\right\}\\
&\ge (b-a)^{-1}\left\{a\int_a^b (b-y)f(y)\,dy + b\int_a^b (y-a)f(y)\,dy\right\}\\
&= \int_a^b yf(y)\,dy.
\end{aligned}$$

- If $a \le y \le x \le b$, then

$$\begin{aligned}
&(b-a)^{-1}\left\{x\int_a^b (b-y)f(y)\,dy + b\int_a^b (y-a)f(y)\,dy\right\}\\
&\ge (b-a)^{-1}\left\{x\int_a^b (b-y)f(y)\,dy + x\int_a^b (y-a)f(y)\,dy\right\}\\
&= \int_a^b xf(y)\,dy.
\end{aligned}$$

- If $a \le y \le b \le x$, then

$$\begin{aligned}
(b-a)^{-1}\left\{x\int_a^b (b-y)f(y)\,dy + x\int_a^b (y-a)f(y)\,dy\right\} &= (b-a)^{-1}\int_a^b (bx - xy + xy - ax)f(y)\,dy\\
&= \int_a^b xf(y)\,dy.
\end{aligned}$$

Hence (6) is satisfied.

Thus

$$\begin{aligned}
E(\max\{X,Y\}) &= \int_{-\infty}^a \max\{x,y\}f(y)\,dy + \int_b^\infty \max\{x,y\}f(y)\,dy + \int_a^b \max\{x,y\}f(y)\,dy\\
&\le \int_{-\infty}^a \max\{x,y\}f(y)\,dy + \int_b^\infty \max\{x,y\}f(y)\,dy\\
&+ (b-a)^{-1}\left\{E(\max\{X,a\})\int_a^b (b-y)f(y)\,dy + E(\max\{X,b\})\int_a^b (y-a)f(y)\,dy\right\}\\
&= E(\max\{X,Y_a^b\}),
\end{aligned}$$

as required.

B.4 Proof of Lemma 5.1:

Let the non-negative random variable X have pdf $f(x)$ and cdf $F(x)$. Since $1 - F(x) = P(X \geq x) = \int_x^\infty f(t)dt$, then

$$\int_0^\infty (1 - F(x))\, dx = \int_0^\infty P(X \geq x)\, dx = \int_0^\infty \int_x^\infty f(t)\, dt dx.$$

Changing the order of integration, noting that $0 \leq x \leq t < \infty$,

$$\int_0^\infty \int_x^\infty f(t)\, dt dx = \int_0^\infty \int_0^t f(t)\, dx dt = \int_0^\infty [x f(t)]_0^t\, dt = \int_0^\infty t f(t)\, dt.$$

Since t is a dummy variable, we may replace it with x and arrive at the desired result:

$$E(X) = \int_0^\infty (1 - F(x))\, dx.$$