

A More General Pandora Rule?

Wojciech Olszewski^a, Richard Weber^b

^a*Department of Economics, Northwestern University, Evanston, IL 60208, U.S.A.*

^b*Statistical Laboratory, University of Cambridge, Cambridge CB3 0WB, U.K.*

Abstract

In a model introduced by Weitzman an agent called Pandora opens boxes sequentially, in whatever order she likes, discovers prizes within, and optimally stops. Her aim is to maximize the expected value of the greatest discovered prize, minus the costs of opening the boxes. The solution, using the so-called Pandora rule, is attractive and has many applications. However, it does not address applications in which the payoff depends on all discovered prizes, rather than just the best of them, nor is it easy to say whether or not some generalized Pandora rule might do so. Here, we establish a sense in which it cannot. We discover that if a generalized Pandora rule is to be optimal for some more general utility, and all model parameters, then the problem can be solved via a second problem having Weitzman's form of utility.

JEL classification: C44, D83

Keywords: Optimal scheduling and stopping; Weitzman's Pandora rule

1. Introduction

In a classic problem first analyzed and solved by Weitzman (1979) an agent called Pandora is presented with n boxes, each of which contains a prize. Pandora can, by paying a known cost c_i , open box i to reveal its prize. The nonnegative value of the prize, denoted x_i^o , is not known until the box is opened, but it has a distribution F_i that is known ex ante.¹ Pandora wishes to choose the order of opening the boxes, and when to stop opening, so as to maximize the expected value the greatest discovered prize net of the sum of the costs of opening boxes. Weitzman's problem is attractive for two reasons. Firstly, it has an enormous number of applications, such as to searching for a house, job, or research project to conduct.

Secondly, the solution is remarkably simple and attractive. Assign to any unopened box, say box i , a *reservation value* (or *reservation prize*), of

$$x_i^\dagger = \inf \{y : y \geq -c_i + E \max[y, x_i^o]\}, \quad (1)$$

*Corresponding author: Richard Weber, email: rrw1@cam.ac.uk, telephone/fax: +44-1223-337944/337956

¹The superscript 'o' is provided as a mnemonic for 'opened' or 'observed'.

the expectation being taken over x_i^o . This is the smallest prize value y for which the agent can do as well by taking y , as by opening box i and taking its prize or y . In (1) and throughout this note we are concerned only with Weitzman’s problem in which costs and reward are undiscounted.

The so-called *Pandora rule*, which is optimal for Weitzman’s problem, is: *open boxes in descending order of reservation values until a prize is found whose value weakly exceeds the reservation value of any unopened box.*

Attractive as it is, Weitzman’s model does not cover an important and large class of problems in which the agent’s utility is not only a function of the prize the agent takes when she stops, but of all the prizes uncovered. For example, a student may benefit from courses she takes while searching for the subject to choose as major; or a person may obtain a flow of utility by dating different partners while looking for a spouse; or an institution may be affected by different forms of operation with which it temporarily experiments before adopting a permanent one.

Weitzman expected that Pandora’s rule would not generalize to such problems. He wrote: “*If some fraction of its reward can be collected from a research project before the sequential search procedure as a whole is terminated, that could negate Pandora’s rule in extreme cases.*” However, Weitzman gave no supporting detailed analysis, and it turns out to be difficult to say whether or not some interesting generalization might be possible. In this note, we fill this gap in understanding by starting with a very general utility and propose a generalized Pandora rule. Then we take as “extreme cases” the requirement that this Pandora rule should be optimal for all choices of costs and prize value distributions, our motivation being that this is true in Weitzman’s problem. We discover that if a generalized Pandora rule is to be optimal under this requirement then the problem can also be solved by solving a second problem having Weitzman’s utility, using his Pandora rule.

2. Model

2.1. The generalized Pandora problem

An agent called Pandora is presented with n boxes, each of which contains a prize. Pandora can, by paying a known cost c_i , open box i to reveal its prize. The nonnegative value of the prize, denoted x_i^o , is not known until the box is opened, but it has a distribution function F_i which is known ex ante.

If S is the set of opened boxes at the point the agent stops, and the vector of the prize values found is $x_S^o = (x_i^o, i \in S)$, then the agent obtains a reward $u(x_S^o)$, expressed as a utility that depends on all the prizes discovered. Pandora’s aim is to maximize the expected value of

$$u(x_S^o) - \sum_{i \in S} c_i. \tag{2}$$

2.2. Assumptions

In undiscounted Weitzman’s problem, $u(x_S^o) = \max_{i \in S} x_i^o$. This utility has the following four special properties. It is *continuous* in its arguments, *symmetric* (in the sense that for any k -tuple of arguments the value is unchanged by permutation of the arguments), *monotone* (in the sense that it is monotone coordinate-wise),

and *submodular*, in the sense that the increase in $u(x)$ obtained by increasing one coordinate of x becomes no greater as any other coordinate becomes greater. That is, for any x_S^o , $\underline{x}_1 < \bar{x}_1$ and $\underline{x}_2 < \bar{x}_2$,

$$u(x_S^o, \bar{x}_1, \bar{x}_2) - u(x_S^o, \underline{x}_1, \bar{x}_2) \leq u(x_S^o, \bar{x}_1, \underline{x}_2) - u(x_S^o, \underline{x}_1, \underline{x}_2). \quad (3)$$

These properties are immediate in the case of $u(x_S^o) = \max_{i \in S} x_i^o$, except submodularity. Submodularity can be verified by considering separately the following two cases: (1) $\underline{x}_2, \bar{x}_2 \leq \max\{\underline{x}_1, x_i^o, i \in S\}$ or $\max\{\bar{x}_1, x_i^o, i \in S\} \leq \underline{x}_2, \bar{x}_2$; (2) $\max\{\underline{x}_1, x_i^o, i \in S\} < \underline{x}_2 < \max\{\bar{x}_1, x_i^o, i \in S\}$ or $\max\{\underline{x}_1, x_i^o, i \in S\} < \bar{x}_2 < \max\{\bar{x}_1, x_i^o, i \in S\}$. In the former case (3) holds with equality, and in the latter case (3) holds with strict inequality.

We consider more generally, problems in which the following assumption is satisfied.

Assumption 1: Rewards are given by a utility function

$$u : \{\emptyset\} \cup \bigcup_{k=1}^{n+1} \mathbb{R}_+^k \rightarrow \mathbb{R}_+,$$

where

- (a) $u(\emptyset) = 0$ and $u(0, x_2, \dots, x_k) = u(x_2, \dots, x_k)$;
- (b) u is continuous, symmetric, nondecreasing and submodular, in the sense that for every $k = 1, \dots, n+1$, the restriction of u to \mathbb{R}_+^k is continuous, symmetric, monotone and submodular.

Some comments on Assumption 1 seem helpful. Part (a) says that prize value 0 can be interpreted as “no prize”. Symmetry means that u depends only on the set of prize values uncovered, not on the order in which they are uncovered. Continuity and monotonicity are standard assumptions. Submodularity is most restrictive, but it is satisfied in many applications, including the Weitzman’s problem and all applications mentioned in Introduction.

3. Result

3.1. The generalized Pandora rule

Suppose a set of boxes $S \subset N = \{1, \dots, n\}$ has been opened, and $i \notin S$. We might ask, what is the smallest prize value whose addition to the set of those already discovered would make it as good to stop as to open box i and then stop? We define generalized reservation value as

$$x_i^*(x_S^o) = \inf\{y : u(x_S^o, y) \geq -c_i + Eu(x_S^o, y, x_i^o), y \geq 0\}, \quad (4)$$

where the expectation is taken over x_i^o , and with the understanding that inf over an empty set is interpreted as $x_i^*(x_S^o) = \infty$. The calculation in (4) makes sense because x_S^o is a vector of length no more than $n-1$ and we

have assumed that u maps a vector of any length no more than $n + 1$ to a real value.^{2,3}

With no loss of clarity, we mostly drop the argument and write the reservation value as x_i^* rather than $x_i^*(x_S^o)$. The generalized Pandora rule is now this: *Open an unopened box with the greatest reservation value, until there is no unopened box whose reservation value strictly exceeds 0.*

In the special case of Weitzman’s problem (that is, when $u(x_S^o) = \max_{i \in S} x_i^o$) his Pandora rule and our generalized Pandora rule coincide. To show this, notice that for any variable x_i , one of the following two cases holds: (1) the inequality

$$u(x_S^o, y) \geq -c_i + Eu(x_S^o, y, x_i^o) \tag{5}$$

is satisfied by some $y \leq \max_{i \in S} x_i^o$; or (2) the inequality may be satisfied only by $y > \max_{i \in S} x_i^o$. In the former case, the inequality is also satisfied by $y = 0$, so $x_i^*(x_S^o) = 0$, and in the latter case $x_i^*(x_S^o) > 0$. (In the latter case, inequality (5) cannot be satisfied by y ’s arbitrarily close to 0, since then it would also be satisfied by $y = 0$ by continuity.) If case (1) applies to all variables x_i , then $x_i^\dagger \leq \max_{i \in S} x_i^o$ and $x_i^*(x_S^o) = 0$ for all uncovered i ’s, and so both Weitzman’s Pandora rule and the generalized Pandora rule stop. If case (2) applies to some uncovered variable x_i , then the rules open the box with the greatest x_i^\dagger or $x_i^*(x_S^o)$, respectively, and it follows directly from the definitions that $x_i^\dagger = x_i^*(x_S^o)$ for the i ’s to which case (2) applies.

3.2. A constraining result on the Pandora rule

We can now state our result. Its proof is in Appendix A.

Theorem 1. *Suppose utility u satisfies Assumption 1, and the generalized Pandora rule maximizes expected value for all costs c_i and distributions F_i . Then for any given c_i and $x_i \sim F_i$ the set of solutions found by the generalized Pandora rule coincides with the set of solutions found by Weitzman’s Pandora rule for some other \bar{c}_i , $\bar{x}_i \sim \bar{F}_i$, and $\bar{u}(x_S^o) = \max_{i \in S} x_i^o$.*

Remark 1. In proving Theorem 1 we allow costs and distributions for which the reservation values are infinite. We believe Theorem 1 is true if costs and distributions are restricted to those for which all reservation values are finite, but the proof is likely to be long and provide no interesting additional insight. See footnote 5 in Appendix A.

We have shown in this note that there do not exist any simple conditions on model parameters for which a non-trivial extension of Weitzman’s result is possible. However, in a working paper Olszewski and Weber (2014)

²One might wonder if the reservation value could be defined as some other function of the cost of opening a box, the distribution prize value inside, and prize values already discovered. We conjecture that if it is always optimal to open boxes in descending order of reservation values then the reservation values must coincide with those described in §3.1, up to a monotone rescaling.

³One might try to generalize further by allowing utility to depend on prizes in closed boxes. For example, Klabjan et al. (2014) study a searcher who must decide whether to accept or reject an object which is characterized by multiple attributes. Before making the decision the searcher can, at some cost, discover the values of selected attributes. For this more ambitious and difficult problem Klabjan et al. were only able to characterize the optimal strategy for very specific classes of utilities, distributions and costs.

we prove that the generalized Pandora rule is optimal if there hold Assumptions 1 and 2 (defined in proof below), and the ordering of the generalized reservation values of covered prizes is independent of both the number of prizes that have already been uncovered and their values. These conditions are met in Weitzman’s model, and in at least one further example for which neither Weitzman’s results or the Gittins index theorem can provide the solution. (We also explain in Olszewski and Weber (2014) how Weitzman’s result can be deduced from the Gittins index theorem for bandit processes; see Gittins and Jones (1974)). Under these conditions the proof of optimality of generalized Pandora rule is engagingly simple.

References

- [1] J. C. Gittins, D. M. Jones (1974), A dynamic allocation index for the sequential design of experiments, in: J. Gani (Ed.), Progress in Statistics, North-Holland, Amsterdam, NL, 1974, pp. 241–66, read at the 1972 European Meeting of Statisticians, Budapest.
- [2] D. Klabjan, W. Olszewski, A. Wolinsky (2014), Attributes, Games and Economic Behavior 88 (6), 190–206.
- [3] W. Olszewski, R. R. Weber (2014), A more general Pandora’s rule?, arXiv:1509.09095 [math.OC].
- [4] M. L. Weitzman (1979), Optimal search for the best alternative, Econometrica 47 (3) 641–654.

Appendix A. Proof of Theorem 1

For the proof, it will be convenient to introduce another property of utility function u , which we label as Assumption 2.

Assumption 2: The marginal benefit of increasing a coordinate of x from 0 to some positive value x_j^o is independent of the values of coordinates of x which are greater than x_j^o . That is, for x_S^o and any $x_j^o \leq \underline{x}_k < \bar{x}_k$, with $j, k \notin S$ and $j \neq k$,

$$u(x_S^o, \bar{x}_k, x_j^o) - u(x_S^o, \bar{x}_k, 0) = u(x_S^o, \underline{x}_k, x_j^o) - u(x_S^o, \underline{x}_k, 0). \quad (\text{A.1})$$

Lemma 1. *Suppose that utility u satisfies Assumption 1 and the (generalized) Pandora rule maximizes expected value for all costs c_i and distributions F_i . Then u satisfies Assumption 2.*

Proof. Consider an arbitrary S , x_S^o , and $j, k \notin S$, with $j \neq k$, and numbers $x_j^o \leq \underline{x}_k < \bar{x}_k$. Suppose there were a violation of (A.1) of the form

$$u(x_S^o, x_j^o, \underline{x}_k) - u(x_S^o, \underline{x}_k) > u(x_S^o, x_j^o, \bar{x}_k) - u(x_S^o, \bar{x}_k).$$

Notice that because of Assumption 1 (submodularity) we cannot have the opposite strict inequality. Applying Assumption 1, we can increase x_j^o to \underline{x}_k and the same inequality will hold. So there exists $\epsilon > 0$ such that

$$u(x_S^o, \underline{x}_k, \underline{x}_k) - u(x_S^o, \underline{x}_k) > u(x_S^o, \underline{x}_k, \bar{x}_k) - u(x_S^o, \bar{x}_k) + \epsilon. \quad (\text{A.2})$$

We now show that if (A.2) is true then the Pandora rule cannot be optimal for all $(c_i, F_i, i \in N)$. To this end, suppose F_j and F_k are degenerate, with $x_j^o = \underline{x}_k^o$ and $x_k^o = \bar{x}_k^o$ with probability 1. Let

$$c_j = u(x_S^o, \underline{x}_k^o, \underline{x}_k^o) - u(x_S^o, \underline{x}_k^o) - \epsilon, \quad c_k = u(x_S^o, \underline{x}_k^o, \bar{x}_k^o) - u(x_S^o, \underline{x}_k^o).$$

The reservation price of x_j is the least nonnegative y such that

$$c_j = u(x_S^o, \underline{x}_k^o, \underline{x}_k^o) - u(x_S^o, \underline{x}_k^o) - \epsilon \geq u(x_S^o, \underline{x}_k^o, y) - u(x_S^o, y). \quad (\text{A.3})$$

Since (A.3) is false for $y = \underline{x}_k^o$ we must have $x_j^* > \underline{x}_k^o$.

The reservation price of x_k is the least nonnegative y such that

$$c_k = u(x_S^o, \underline{x}_k^o, \bar{x}_k^o) - u(x_S^o, \underline{x}_k^o) \geq u(x_S^o, \bar{x}_k^o, y) - u(x_S^o, y). \quad (\text{A.4})$$

Suppose y is such that (A.3) holds, and therefore $y \geq x_j^* > \underline{x}_k^o$. Then (A.4) also holds, since

$$\begin{aligned} & u(x_S^o, \bar{x}_k^o, y) - u(x_S^o, y) \\ &= u(x_S^o, \bar{x}_k^o, y) - u(x_S^o, \underline{x}_k^o, y) + u(x_S^o, \underline{x}_k^o, y) - u(x_S^o, y) \\ &\leq u(x_S^o, \bar{x}_k^o, y) - u(x_S^o, \underline{x}_k^o, y) + u(x_S^o, \underline{x}_k^o, \underline{x}_k^o) - u(x_S^o, \underline{x}_k^o) - \epsilon \\ &= c_k + [u(x_S^o, \bar{x}_k^o, y) - u(x_S^o, \underline{x}_k^o, y)] - [u(x_S^o, \underline{x}_k^o, \bar{x}_k^o) - u(x_S^o, \underline{x}_k^o, \underline{x}_k^o)] - \epsilon \\ &\leq c_k, \end{aligned}$$

where the first inequality is by (A.3) and the second inequality follows by using Assumption 1 (submodularity) to see that, since $y \geq \underline{x}_k^o$, the first square-bracketed term is no greater than the second.

From this it follows that $x_k^* \leq x_j^*$. Thus, according to the Pandora rule, it would be optimal to begin by next uncovering x_j . However, the payoff obtained by uncovering x_k first and then stopping is strictly greater than the payoff obtained by uncovering x_j first and then stopping if

$$u(x_S^o, \bar{x}_k^o) - c_k > u(x_S^o, \underline{x}_k^o) - c_j. \quad (\text{A.5})$$

On substituting for c_j and c_k we find that (A.5) is the same as (A.2).

The payoff obtained by uncovering x_k first and then stopping is also strictly greater than the payoff of uncovering both if

$$u(x_S^o, \bar{x}_k^o) - c_k > u(x_S^o, \underline{x}_k^o, \bar{x}_k^o) - c_j - c_k. \quad (\text{A.6})$$

On substituting for c_j we find that (A.6) is also the same as (A.2).

Now suppose that apart from x_j and x_k , all other prizes should certainly stay covered. For example, each other value might be $x_i^o = 0$ with probability 1, and $c_i > 0$. We have argued that uncovering x_k first and then stopping is strictly better than uncovering x_j first and then either stopping or uncovering x_k . As $x_k^* \leq x_j^* > 0$, the Pandora rule dictates that it is optimal to uncover prize x_j first. As this is false, we must conclude that if the Pandora rule is optimal then (A.2) must be false, and thus Lemma 1 is true. \square

Lemma 2. *Suppose that utility u satisfies Assumptions 1 and 2. For any x_S^o , let x_ℓ denote the ℓ -th greatest coordinate of x_S^o .⁴ Then,*

(a) *there exist functions $g_\ell : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, $\ell = 1, 2, \dots, n$, such that for any x_S^o we have*

$$u(x_S^o) = \sum_{\ell=1}^{|S|} g_\ell(x_\ell),$$

(b) *$g_\ell(x)$ is (weakly) increasing in x and (weakly) decreasing in ℓ ,*

(c) *$g_\ell(x) - g_{\ell+1}(x)$ is weakly increasing in x .*

(d) *there exists some g such that $g_2 = \dots = g_n = g$, and hence*

$$u(x_S^o) = u(\max_{i \in S} x_i^o) - g(\max_{i \in S} x_i^o) + \sum_{i \in S} g(x_i^o) \quad (\text{A.7})$$

where $u - g$ is nondecreasing, and $u(0) = g(0) = 0$.

Once we know (A.7) is true, then we can complete a proof of Theorem 1 as follows. Consider Weitzman's problem with cost \bar{c}_i , and prize values, \bar{x}_i^o defined by

$$\bar{c}_i = c_i - E g(x_i^o),$$

$$\bar{x}_i^o = u(x_i^o) - g(x_i^o).$$

The payoff obtained as $E[-\sum_{i \in S} \bar{c}_i + \max_{i \in S} \bar{x}_i^o]$ will be the same as obtained with (A.7). This completes the proof of Theorem 1. It remains to prove the Lemma 2.

Proof. To prove (a), define

$$g_\ell(x) = u(\underbrace{x, \dots, x}_{\ell \text{ times}}) - u(\underbrace{x, \dots, x}_{\ell-1 \text{ times}}).$$

Then for $S = \{1, \dots, k\}$ and $x_S^o = (x_i, i \in S)$,

$$\begin{aligned} u(x_S^o) &= u(x_1, \dots, x_k) \\ &= u(x_1, \dots, x_k) - u(x_1, \dots, x_{k-1}, 0) + u(x_1, \dots, x_{k-1}) \end{aligned} \quad (\text{A.8})$$

$$= u(\underbrace{x_k, \dots, x_k}_{k \text{ times}}) - u(\underbrace{x_k, \dots, x_k, 0}_{k-1 \text{ times}}) + u(x_1, \dots, x_{k-1}) \quad (\text{A.9})$$

$$= g_k(x_k) + u(x_1, \dots, x_{k-1})$$

$$= \sum_{\ell=1}^k g_\ell(x_\ell),$$

where (A.9) follows from (A.8) by repeated application of Lemma 1.

⁴Of course, x_ℓ is a function of x_S^o , but we drop the argument with no loss of clarity.

For (b), the fact that $g_\ell(x)$ is a decreasing function of ℓ follows from Assumption 1 (submodularity). The fact that $g_\ell(x)$ is increasing in x can be seen by taking $x < x'$, and observing that

$$g_\ell(x) = u(\underbrace{x, \dots, x'}_{\ell-1 \text{ times}}) - u(\underbrace{x', \dots, x'}_{\ell-1 \text{ times}}) \leq g_\ell(x'),$$

where the equality is by Lemma 1 and the inequality is by Assumption 1.

For (c), we note that if $x < x'$,

$$\begin{aligned} g_\ell(x) - g_{\ell+1}(x) &= [u(\underbrace{x, \dots, x}_{\ell \text{ times}}) - u(\underbrace{x, \dots, x}_{\ell-1 \text{ times}})] - [u(\underbrace{x, \dots, x}_{\ell+1 \text{ times}}) - u(\underbrace{x, \dots, x}_{\ell \text{ times}})] \\ &= [u(\underbrace{x', \dots, x', x}_{\ell-1 \text{ times}}) - u(\underbrace{x', \dots, x'}_{\ell-1 \text{ times}})] - [u(\underbrace{x', \dots, x', x}_{\ell \text{ times}}) - u(\underbrace{x', \dots, x'}_{\ell \text{ times}})] \\ &\leq [u(\underbrace{x', \dots, x'}_{\ell \text{ times}}) - u(\underbrace{x', \dots, x'}_{\ell-1 \text{ times}})] - [u(\underbrace{x', \dots, x'}_{\ell+1 \text{ times}}) - u(\underbrace{x', \dots, x'}_{\ell \text{ times}})] \\ &= g_\ell(x') - g_{\ell+1}(x'). \end{aligned}$$

The second line follows by Lemma 1, and the third line by Assumption 1 (submodularity).

For (d) We prove $g_2 = g_3$. The proof of $g_\ell = g_{\ell+1}$, $\ell > 2$ follows by examining an instance in which the first $\ell - 2$ prizes uncovered are ones with $c_i = 0$ (and reservation values ∞) and their uncovered values are greater than any values that can be found amongst the prizes which remain uncovered at that point.

(i) Assume first that g_3 is not equal to 0. Then there exists x_0 such that $g_1(x_0) \geq g_2(x_0) > g_3(x_0) > 0$. Consider three variables, x_1 , x_2 and x_3 with the same degenerate distribution, having $x_i^o = x_0$, with probability 1, $i = 1, 2, 3$. Let costs be chosen so

$$c_3 = 0 \leq c_1 < g_3(x_0) < c_2 < g_2(x_0). \quad (\text{A.10})$$

We proceed to show the Pandora rule cannot be optimal.

Firstly, it follows from (A.10) that for all y we have $u(y) < -c_i + Eu(y, x_i^o)$. Hence initially, when $S = \emptyset$, all three prizes have reservation value ∞ . So if the Pandora rule is optimal then it must be optimal to uncover any of them first. Suppose x_2 is uncovered first, and then x_3 (which still has reservation value ∞).⁵ It is now strictly best to uncover x_1 if and only if

$$g_1(x_0) + g_2(x_0) < -c_1 + g_1(x_0) + g_2(x_0) + g_3(x_0)$$

which is true because $c_1 < g_3(x_0)$. The payoff is that of uncovering all three prizes.

⁵The proof of Lemma 2 (d) is where using variables with infinite reservation values facilitates analysis. Variables x_1, x_2, x_3 with finite reservation values would require non-degenerate distributions, and controlling for the desired orderings of reservation values (initial and after uncovering one of them) would require an elaborate construction of these variables.

Alternatively, if we uncover x_1 first, followed by x_3 , it is now strictly best not to uncover x_2 , since $c_2 > g_3(x_0)$. The difference in expected payoffs of the strategy which uncovers x_1, x_3 and of that which uncovers x_2, x_3, x_1 is

$$[-c_1 - c_3 + g_1(x_0) + g_2(x_0)] - [-c_1 - c_3 - c_2 + g_1(x_0) + g_2(x_0) + g_3(x_0)] = c_2 - g_3(x_0),$$

which is positive, whereas if the Pandora rule were optimal this difference should be no greater than 0.

(ii) Now consider the special case in which $g_3 = 0$ and x_0 is such that $g_1(x_0) \geq g_2(x_0) > g_3(x_0) = 0$. Suppose x_i is a variable such that x_i^o is equal to 0 or x_0 with probabilities $q_i = 1 - p_i$ and p_i respectively, and where $c_i/p_i < g_2(x_0)$. Consider the class of prizes like this, for varying p_i and c_i . All have initial reservation value ∞ . Suppose a prize in this class is uncovered and reveals value x_0 . Subsequent to this, the reservation value of another variable in the class is now the least y , with $y \leq x_0$, such that

$$g_1(x_0) + g_2(y) \geq -c_i + g_1(x_0) + p_i g_2(x_0) + (1 - p_i)g_2(y),$$

i.e. the least y such that $g_2(y) \geq -c_i/p_i + g_2(x_0)$. Since $c_i/p_i < g_2(x_0)$, the reservation value is positive.

So suppose we start with three prizes in this class. We uncover one and it takes value x_0 . The other two prizes have now positive reservation values, the greater of which is for the prize with least value of c_i/p_i . In following the Pandora rule we may start by uncovering any prize initially, and then continue by uncovering prizes in increasing order of c_ℓ/p_ℓ , until either two values of x_0 are revealed or all three prizes are uncovered.

If we uncover the prizes in the order x_i, x_j, x_k then the expected payoff is

$$\begin{aligned} & -c_i + p_i[g_1(x_0) - c_j + p_j g_2(x_0) + q_j[-c_k + p_k g_2(x_0)]] \\ & + q_i[-c_j + p_j[g_1(x_0) - c_k + p_k g_2(x_0)] + q_j[-c_k + p_k g_1(x_0)]] \\ & = (c_k/p_k)p_i p_j p_k + \sigma, \end{aligned}$$

where σ is an expression that is symmetric in i, j, k . So if $c_i/p_i < c_j/p_j < c_k/p_k < g_2(x_0)$ then it is strictly better to begin by uncovering x_i or x_j , than to begin by uncovering x_k . Thus optimality of the Pandora rule is incompatible with $g_2(x_0) > g_3(x_0) = 0$. \square