1: INTRODUCTION TO PROOFS

STEVEN HEILMAN

1. Introduction

By the end of this course, we would like to be able to understand and create proofs. Since logic is the foundation of a mathematical proof, we begin the course with basic logic.

However, even before we discuss logic, we have to discuss precision of language. Precision of language is necessary in creating logical statements. To give an example of how logic and precision of language are important, I want to briefly discuss the following paraphrased sentence that I found in a news article:

> "MP3 files only have five percent of the sounds of an original audio recording."

We first note that the notion of "sound" is a bit vague, and perhaps it is difficult to measure. The language that we typically write and speak naturally has this vagueness, so maybe we are not concerned by this vagueness. Yet, in this course, we need to make more precise sentences. We ultimately want to prove things to be true, and in order to do that we need to express incontrovertible assertions. The intentions of the author of this sentence are not entirely clear, but I think the author meant to say:

> "An MP3 file only has five percent of the file size of its corresponding WAV file."

This sentence is more precise than the previous sentence. In order to write mathematics well, we need to transform our writing in this way. We all know from our experience with MP3 files that this sentence is often true, though perhaps ten percent is more accurate. So, from empirical reasoning, the following is a true sentence:

> "Ninety-five percent of MP3 files have less than ten percent of the file size of their corresponding WAV files."

This statement still does not have the mathematical quality that we want. Empirical reasoning does not always lead to a true statement. However, empirical reasoning can often help us to conjecture something that is true. So let us begin to build up statements at the most basic level.

As an aside, I think the author of our cited sentence made the following assumption: a decrease to five percent of the file size must mean that the listener only hears five percent of the "sound" of the original music. This assumption is actually false. What do I mean by this? Well, first of all, we all know that our MP3s sound just fine. More to the point, there is a rigorous mathematical theory that says that "most" of the sound of the compressed MP3 is preserved. Yes, you can actually prove that MP3s work well. And without these proofs, MP3s would not exist. However, formulating precise statements like these requires background beyond this course.

Date: February 29, 2012.

2. Truth Tables, De Morgan, and Nash

In class, the statements P and Q have been discussed. When we have statements P and Q, we want to perform some basic operations on them. The following table provides a nice summary of these operations. Recall: $\sim$ means "not", $\wedge$ means "and", $\vee$ means "or", $P \Rightarrow Q$ means "if P then Q," and $P \Leftrightarrow Q$ means $(P \Rightarrow Q) \wedge (Q \Rightarrow P)$.

P	Q	$\sim P$	$\sim Q$	$P \wedge Q$	$P \vee Q$	$P \Rightarrow Q$	$P \Leftrightarrow Q$	$(\sim P) \vee (\sim Q)$	$\sim (P \wedge Q)$
T	T	F	F	T	T	T	T	F	F
T	F	F	T	F	T	F	F	T	T
F	T	T	F	F	T	T	F	T	T
F	F	T	T	F	F	T	T	T	T

Table 1. A Truth Table

One reads the table as follows. The first two entries on the left side of a row define whether or not P and Q are true. For example, in the fourth row from the top, P is made to be false, and Q is made to be true. Within this row, operations on P and Q are calculated. For example, in the fourth row from the top, we see that $P \wedge Q$ is false, assuming P is false and Q is true, and so on.

Definition 2.1. The expression $P \Rightarrow Q$ is defined to be the statement $(\sim P) \vee Q$.

Exercise 2.2. Check that this definition of $P \Rightarrow Q$ agrees with the entries of the truth table above.

Using the truth table, we can complete an exercise from the first homework. The following is an example of a proof by exhaustion. We will make a finite list of cases, and then prove something about each separate case.

Proposition 2.3. ***(De Morgan's Law)*** *Let A and B be sets in some universe X. Then $(A \cap B)^c = A^c \cup B^c$*

Proof. Let $x \in X$. For the set A, we have a dichotomy: either $x \in A$ or $x \notin A$. Similarly, for B, we have a dichotomy: either $x \in B$ or $x \notin B$. Since we have two sets and two possibilities, there are exactly $2^2 = 4$ mutually exclusive cases to consider for the location of x: (i) $x \in A$ and $x \in B$ (ii) $x \in A$ and $x \notin B$ (iii) $x \notin A$ and $x \in B$ (iv) $x \notin A$ and $x \notin B$. Let P be the event $x \in A$ and let Q be the event $x \in B$. We can then translate the four cases above into statements in P and Q: (i) $P \wedge Q$ (ii) $P \wedge (\sim Q)$ (iii) $(\sim P) \wedge Q$ (iv) $(\sim P) \wedge (\sim Q)$. Note also that $x \in (A \cap B)^c$ if and only if $\sim (P \wedge Q)$ is true. Also, $x \in A^c \cup B^c$ if and only if $(\sim P) \vee (\sim Q)$ is true. We want to show that $x \in (A \cup B)^c$ if and only if $x \in A^c \cap B^c$. So, to prove this, it suffices to show that $(\sim (P \wedge Q))$ is true if and only if $(\sim P) \vee (\sim Q)$ is true.

The Proposition will then be proven if: given any of the four cases (i),(ii),(iii),(iv), we determine that the truth of $\sim (P \wedge Q)$ is equal to the truth of $(\sim P) \vee (\sim Q)$. Now, cases (i) through (iv) correspond exactly to the assumptions of the last four rows of Table 1. Moreover, in Table 1, the column of $\sim (P \wedge Q)$ is equal to the column of $(\sim P) \vee (\sim Q)$. We conclude that $(\sim (P \wedge Q))$ is true if and only if $(\sim P) \vee (\sim Q)$ is true. The Proposition is therefore proven. □

Exercise 2.4. Explain why this proof is the same as the "usual" proof by picture that $(A \cap B)^c = A^c \cup B^c$.

We will now give a similar proof of a completely different fact. We will describe the problem known as the **Prisoner's Dilemma**. You may recall a similar situation from the movie The Dark Knight. Suppose you and a friend have been arrested for some wrongdoing. The police place each of you in separate rooms, and you are unable to communicate. The police offer the following to you, and you must make a decision before you leave the room. If you will testify against your friend, and your friend will stay silent, then you will go free, and your friend spends five years in jail. If your friend decides to testify against you, and you decide to stay silent, then your friend will go free, and *you* spend five years in jail. If you both decide to testify against each other, then you both spend three years in jail. If you both decide that you will stay silent, you will both spend one year in jail. (The police give your friend an identical offer.) You want to minimize the amount of time that you spend in jail. What should you do?

Definition 2.5. In the Prisoner's Dilemma, suppose you and your friend have made your decisions about whether or not you will confess. We say that your strategies are in a **Nash equilibrium** if the following holds. With your friend's decision considered fixed, you cannot gain anything by changing your decision. And with your decision considered fixed, your friend cannot gain anything by changing her decision.

Theorem 2.6. ***(Prisoner's Dilemma)*** *There is exactly one Nash equilibrium for the Prisoner's Dilemma. In this equilibrium, both you and your friend testify.*

Proof. We analyze each of the four possible strategies.

Strategy 1: You testify and your friend does not. If you are going to testify, then your friend will spend two years less time in jail if she also decides to testify. Therefore, in the case of Strategy 1, it is better for her to change her decision. So Strategy 1 is not in equilibrium.

Strategy 2: You stay silent and your friend testifies. If your friend is going to testify, then you will spend two years less time in jail if you also decide to testify. Therefore, in the case of Strategy 2, it is better for you to change your decision. So Strategy 2 is not in equilibrium.

Strategy 3: You both stay silent. If your friend is going to stay silent, then you will spend one year less in jail if you testify against your friend. Therefore, in the case of Strategy 3, it is better for you to change your decision. So Strategy 3 is not in equilibrium.

Strategy 4: You both testify. If your friend is going to testify, it is best for you to also testify. (If you do not testify, you will spend two more years in jail). Similarly, if you are going to testify, then it is best for your friend to also testify. (If she does not testify, then she will spend two more years in jail). Therefore, Strategy 4 is in a Nash equilibrium. □

Remark 2.7. Situations such as the Prisoner's Dilemma may arise, e.g. in political or business negotiations. Note that the Nash equilibrium is not necessarily the best outcome for both parties!

3. A Problem with Induction

Before we see more examples of proofs, I want to pose a problem. Sometimes I write a proof of some statement, and I know I might have a mistake somewhere, but I cannot find the mistake. You may find yourself in a similar situation. We will now emulate this situation.

Problem 3.1. The following proof will have a mistake somewhere. Test your understanding of induction by trying to find the mistake.

Claim: All horses on Earth are the same color.

Proof. We prove the claim by induction. Let k be a positive integer. In the case $k = 1$, a single horse has the same color as itself, so the case $k = 1$ of the induction is known. We now assume by induction that each set of k horses is of the same color. We want to show that a set of $k + 1$ horses is of the same color. Suppose I have a set of $k + 1$ horses. If I remove one horse from this set of $k + 1$ horses, I have k horses of the same color. Label this set of k horses as A. All horses in the set A are the same color, by the assumption for sets of k horses.

Now, take the set of $k + 1$ horses and remove a different horse from this set than the one that we removed before. Label this new set of k horses as B. All horses in the set B are the same color, by the assumption for sets of k horses. Since A and B have some horses in common, the $(k + 1)$ horses all must have the same color. We have therefore completed the induction, and the claim is proven. □

Clearly there are horses of different colors. Where did I go wrong?

4. Examples of Proofs

We now present some proofs of certain mathematical statements. These proofs may involve several statements that may take some time to read and understand, but I hope they at least give an idea of what a proof should "look like" and of acceptable mathematical discourse.

As our first example, we give a well known argument by contradiction. We will show that there are infinitely many prime numbers. This statement may seem to be true, but it may not be immediately obvious *how* to prove that the statement is true. In undergraduate mathematics, this predicament is common. Indeed, one of the goals of undergraduate mathematics is exactly to learn how to prove statements correctly, even ones that may at first seem obvious. However, be aware that some statements may seem obvious, though they are difficult to prove, or even incorrect!

The argument below is attributed to Euclid, i.e. it is over 2000 years old. Perhaps it is a testament to the power of mathematics that such a thing could last for so long. Before we begin we recall the definition of a prime number.

Definition 4.1. Let p be a positive integer. We say that p is a **prime** number if $p > 1$, and if we write $p = ab$ for a, b positive integers, then either $a = 1$ or $b = 1$.

Note that, by this definition, 1 is not a prime number. Recall that the first few primes, starting from 2 are $2, 3, 5, 7, 11, 13, 17, \ldots$. The following theorem is well known, so we give its proof as an exercise.

Exercise 4.2. (**The Fundamental Theorem of Arithmetic**) Let $n > 1$ be a positive integer. Then n can be written uniquely as a product of its prime factors. That is, given n a positive integer, there exists a unique positive integer m and there exist unique primes $p_1 < p_2 < \cdots < p_m$ and unique positive integers $a_1, \ldots, a_m$ such that

$$n = (p_1)^{a_1} \cdot (p_2)^{a_2} \cdots (p_m)^{a_m}$$

(Hint: Let $n > 1$ be an arbitrary positive integer. What is the negation of the definition of a prime number? How do we know that the factorization has finite length?)

For example, $60 = 2^2 \cdot 3 \cdot 5$.

Theorem 4.3. ***(Euclid)*** *There are infinitely many prime numbers. Thus, for any positive integer N, there exist more than N prime numbers.*

Proof. We argue by contradiction. Assume that there are only N prime numbers, where N is a positive integer. Label these prime numbers $p_1, p_2, \ldots, p_{N-1}, p_N$. Define p by the formula

$$p = (p_1 \cdot p_2 \cdots p_{N-1} \cdot p_N) + 1 \qquad (*)$$

By the Fundamental Theorem of Arithmetic (Exercise 4.2), we can write $p = p' \cdot n$ where p' is a prime number and n is a positive integer. Since p' is prime, there exists $i \in \{1, \ldots, N\}$ such that $p' = p_i$. Combining this fact with $(*)$, we see that $(p_1 \cdots p_N) + 1 = p_i n$, i.e.

$$p_1 \cdots p_{i-1} \cdot p_{i+1} \cdots p_N - n = \frac{1}{p_i} \qquad (**)$$

Since p_i is a prime, $p_i > 1$. So, the left side of $(**)$ is an integer, while the right side of $(**)$ is not an integer. Since we have arrived at a contradiction, the Theorem is proven. □

In the following Theorem, there will be certain assertions that I will *not* justify rigorously. These assertions can be justified, but doing so would use material outside of this course. Try to find where these gaps in reasoning occur.

Theorem 4.4. ***(Euler's Solution to the Basel Problem)***

$$\sum_{n=1}^{\infty} \frac{1}{n^2} = \frac{\pi^2}{6}$$

Proof. From calculus we know that, for all $x \in \mathbb{R}$ such that $x \neq 0$, the following formula holds

$$\frac{\sin(\pi x)}{\pi x} = \sum_{n=1}^{\infty} \frac{(\pi)^{2n-2} x^{2n-2} (-1)^{n-1}}{(2n-1)!} \qquad (*)$$

Recall that $\sin(\pi x) = 0$ if and only if $x = k$ with $k \in \mathbb{Z}$. Since a polynomial is a product of its zeros, the following infinite product formula holds

$$\begin{aligned}\frac{\sin(\pi x)}{\pi x} &= \cdots (1 - \frac{x}{3})(1 - \frac{x}{2})(1 - x)(1 + x)(1 + \frac{x}{2})(1 + \frac{x}{3}) \cdots \\ &= (1 - x^2)(1 - \frac{x^2}{4})(1 - \frac{x^2}{9}) \cdots\end{aligned}$$

To check that the constant on the left side is correct, note that as x approaches 0, the left side goes to 1 and the right side goes to 1. Multiplying the terms in the infinite product and collecting like terms gives

$$\frac{\sin(\pi x)}{\pi x} = 1 - x^2 \left(1 + \frac{1}{4} + \frac{1}{9} + \cdots\right) + x^4(\cdots) + \cdots \qquad (**)$$

Equating the resulting x^2 terms from $(*)$ and $(**)$ shows that

$$-\frac{\pi^2}{3!} = -\left(1 + \frac{1}{4} + \frac{1}{9} + \cdots\right)$$

The result is therefore proven. □

Remark 4.5. Can you make a similar statement for $\sum_{n=1}^{\infty}(1/n^4)$?

5. Medical Residency Matching Algorithm

Every year, thousands of medical students apply to residency programs. After the application and interview process, each student ranks a few schools in a list (most preferred, second most preferred, etc.) Also, each school ranks their applicants in a list (most preferred, second most preferred, etc.) With these inputs, a computer algorithm decides where all students end up. We will describe this algorithm below.

In a sense to be described later, each student ends up with her best choice. In the actual assignment of residency programs, some students may not be given any program at all. Also, the desire of (real or fake) couples to stay close to each other creates significant complications for the algorithm. We will make some simplifying assumptions to avoid these issues. In particular, we will not take into account the preferences of couples.

Suppose there are n students $\{S_1, \ldots, S_n\}$, and n schools $\{C_1, \ldots, C_n\}$. For simplicity, we assume that every student makes a preference list that includes all n schools, every school makes a preference list that includes all n students, and each school will only accept one student.

A **matching** is defined as a function $f\colon \{S_1, \ldots, S_n\} \to \{C_1, \ldots, C_n\}$ such that, for every integer j with $1 \leq j \leq n$ and for every C_j, there exists an integer i with $1 \leq i \leq n$ such that $f(S_i) = S_j$. For an integer i with $1 \leq i \leq n$, we say that S_i is matched to $f(S_i)$. Note that, by this definition, every student is matched to exactly one school, and every school is matched to exactly one student.

Let S, S' be students and let C, C' be schools. A matching is called **unstable** if the following situation occurs. Student S is matched to C, and student S' is matched to C'. However, S prefers C' over C, and C' prefers S over S'. That is, there exists a pair that is mutually dis-satisfied. A matching is called **stable** if the matching is not unstable. A school C is called a **valid destination** for S if there exists a stable matching such that S is matched to C. The **best destination** for S is the highest ranked school for S, among all valid destinations for S. (Here, when we say rank, we mean the rank that S places on schools within her preference list.)

The following iterative algorithm is used in the residency matching procedure. At each iteration, some student applies to some school, and this student is either accepted or rejected. The process then continues. If the student is accepted by the school, we say that the student is **matched** to the school and that the school is matched to the student. If a student is not currently matched to any program, we say that the student is **unmatched**. Similarly, if a school is not currently matched to any student, we say that the school is unmatched.

Algorithm 5.1. Let S be any unmatched student who has not applied to every school. (If no such student S exists, terminate the algorithm.) Let C be the highest ranked school to which S has not yet applied. (Here, when we say rank, we mean the ranks that S created for schools.)

Case 1: C is not matched to any student. In this case, match S to C.

Case 2: C is matched to a student S'. In this case, if C prefers S' over S, leave S unmatched. If C prefers S over S', match S to C, and make S' unmatched.

One iteration of the algorithm is now complete. Now return to the beginning again.

Note that we had some freedom in the choice of the unmatched student at the beginning of each iteration. We will see later that our choice here will have no effect on the matching that is produced in the end.

Claim 1: The algorithm terminates after at most n^2 iterations.

Claim 2: A school C is initially not matched to any student, and then C is always matched to some student after the first one has applied to C. While C is matched to a student, the student improves in rank or stays the same, at each iteration. (Here, when we say rank, we mean the rank that C places on students.)

Claim 3: Let S be a student. If S is unmatched at some point in the algorithm, then there exists a school C to which she has not yet applied.

Claim 4: When the algorithm terminates, every student is matched to some school. Conversely, every school is matched to exactly one student. So, Algorithm 5.1 produces a matching.

Claim 5: The matching produced by Algorithm 5.1 is stable.

Claim 6: The school that a student S receives at the end of an execution of Algorithm 5.1 is the best destination for S.

Remark 5.2. At the beginning of every iteration of Algorithm 5.1, we have some freedom to choose an unmatched student. However, Claim 6 asserts that the matching that the algorithm produces is the same, regardless of these free choices.

Proof. (of Claim 1) We have n students, and each student can make at most n applications to schools. So, there are at most $n \cdot n = n^2$ distinct applications that students make to schools. So, there are at most n^2 iterations of the algorithm before termination. □

Proof. (of Claim 2) The school C has no student until one student applies to C in Case 1 of the algorithm. Afterwards, a student can only apply to the school C in Case 2 of the algorithm. In Case 2, the rank of the student (from the perspective of C) can only stay the same or increase. Moreover, Case 2 maintains that school C is matched to some student. □

Proof. (of Claim 3) We argue by contradiction. Suppose S is unmatched, and S has applied to all schools. Since S has applied to all schools, Claim 2 implies that all schools have been matched to some student. But there are n schools and less than n students that are matched. We therefore have a contradiction. We conclude that Claim 3 holds. □

Proof. (of Claim 4) We argue by contradiction. Suppose the algorithm ends and either: (i) some school has no student, or (ii) some student has no school. Note that (i) implies that (ii) occurs, since each school is matched to at most one student. So, for the sake of obtaining a contradiction, we may assume that (ii) occurs. That is, we assume that the algorithm ends and some student S is unmatched. By the termination condition of Algorithm 5.1, S must have applied to all schools. But then Claim 3 is contradicted. Since Claim 3 is true, we have obtained a contradiction. We conclude that Claim 4 is true. □

Proof. (of Claim 5) We argue by contradiction. Assume that the algorithm terminated and there exist students S, S' and schools C, C' such that S is matched to C and S' is matched to C'. Moreover, assume that:

$$S \text{ prefers } C' \text{ over } C, \text{ and } C' \text{ prefers } S \text{ over } S'$$

By the definition of Algorithm 5.1, S applied to school C last (with respect to all applications that S made), and S remained matched to C afterwards. Similarly, S' applied to school C' last (with respect to all applications that S' made), and S' remained matched to C' afterwards. We now consider two cases:

$$(i)\, S \text{ applied to } C' \text{ at some point in the algorithm, or } (ii)\, S \text{ never applied to } C'$$

Note that S applied to C, and C has a lower rank than C' (according to S). Moreover, if S is chosen at the beginning of an iterationg of the algorithm, S must always apply to her highest ranked school. Therefore, S must have applied to C', i.e. Case (ii) cannot occur, and we may assume that Case (i) occurs.

In Case (i), since S ends the algorithm matched to C, there must exist a student S'' such that, at some point in the algorithm, S is rejected by C' in favor of S''. That is, C' prefers S'' over S. Since C' ends the algorithm matched to S', Claim 2 says that C' prefers S' over S''. Combining these two preferences of C', we conclude that C' prefers S' over S. However, we initially assumed that C' prefers S over S'. Since we have reached a contradiction, Claim 5 is proven. □

Proof. (of Claim 6) We argue by contradiction. Suppose the algorithm terminates, and suppose that some student is not matched to her best destination. Since a student applies to schools in descending order of preference, the definition of the best destination shows that some student applied to her best destination at some point in the algorithm. Since there are only a finite number of iterations of the algorithm by Claim 1, there exists a first time when a student S is rejected by one of her valid destinations, C. Suppose the algorithm terminates and S ends up matched to C'. When C rejects S, C is either already matched to some student, or at some later time, another student applies to C, causing C to break off a match with S. In either case, the rejection of S is caused by some student S', and C prefers S' over S.

Since C is the best destination of S, there exists some stable matching f in which S is matched to C. In the matching f, suppose S' is matched to some school C'. Recall that the rejection of S by C was chosen as the first such rejection of any student by any valid destination, within some fixed execution of Algorithm 5.1. So, at the point of the rejection of S by C, we know that S' has not yet been rejected by a valid destination. So, at this point in the execution of the algorithm, S' has not applied to C', since C' is a valid destination for S'. Since S' applies to schools in decreasing order of preference (within the execution of Algorithm 5.1), it follows that S' prefers C over C'.

In summary, S' prefers C over C', and C prefers S' over S. However, in the stable matching f, S is matched to C, and S' is matched to C'. Therefore, the matching f is unstable, a contradiction. We therefore conclude that Claim 6 is true. □

6. A Proof from the Field of Optimization

In applications of mathematics, one often has some sort of procedure that one wants to optimize. For example, I have some materials, and I want to construct a building in the shortest amount of time possible, so that my costs are minimized. However, it is not always obvious how to minimize your costs. Moreover, even if the best strategy *seems* obvious, it is not always clear how to *prove* that your strategy is the best. And you can only truly know that your strategy is the best once you have a proof of that fact. Otherwise, you can only speculate. I hope the following example shows that a need for a proof can naturally arise. This proof will be more involved than the ones we have shown above, and it may be difficult to understand upon a first reading. I also hope that this proof shows that there is a great deal more to mathematics than logic and number theory, since the examples that we have dealt with mostly fall into the latter categories. On the contrary, most of mathematics is neither logic nor number theory.

Suppose I have a certain number of bakeries that produce loaves of bread. Also, suppose I have a certain number of restaurants that receive shipments of bread. Each bakery has its own individual daily production, and each restaurant has its own individual daily consumption. (We assume for simplicity that the production and consumption are the same every day.) However, shipping one loaf from a certain bakery B to a certain restaurant R incurs some cost $c(B, R)$. Here $c(B, R)$ is a positive real number. Our goal is to minimize the total cost of transporting the loaves of bread.

Certain questions naturally arise here. Does an optimal transportation plan of loaves exist? And if so, how can we find this optimal plan? The answer to the first question is yes, since there are only a finite number of loaves, a finite number of bakeries, and a finite number of restaurants. So there are only a finite number of possible transportation plans, and at least one of these plans must have a cost that is less than or equal to the other plans. The harder question is, how can we find such a plan?

We begin with an observation. Suppose B_1, B_2, B_3 are bakeries and R_1, R_2, R_3 are restaurants. Assume that B_1 ships a loaf to R_1, B_2 ships a loaf to R_2 and B_3 ships a loaf to R_3. Now, if there is a way to re-route these three loaves to decrease the cost $c(B_1, R_1) + c(B_2, R_2) + c(B_3, R_3)$, then our transport plan is not optimal. For example, if

$$c(B_1, R_2) + c(B_2, R_3) + c(B_3, R_1) < c(B_1, R_1) + c(B_2, R_2) + c(B_3, R_3) \tag{1}$$

Then we can re-route the loaves as depicted in Figure 1, and the total cost will be decreased. Instead of sending a loaf from B_1 to R_1, we send a loaf from B_1 to R_2. Instead of sending a loaf from B_2 to R_2, we send a loaf from B_2 to R_3. And instead of sending a loaf from B_3 to R_3, we send a loaf from B_3 to R_1.

If we have an optimal transportation plan, then we do not want a situation as in (1) to occur. We therefore negate this statement and turn it into a definition. We will expect that our optimal transportation plan will satisfy this definition.

Definition 6.1. Let n, m be positive integers, let $\mathcal{B}$ be the set of bakeries and let $\mathcal{R}$ be the set of restaurants. Let $A \subseteq \mathcal{B} \times \mathcal{R}$. We say that A is **cyclically monotone** if, for any positive integer N and for any set $(B_1, R_1), \ldots, (B_N, R_N)$ of points in A, the following inequality holds

$$\sum_{i=1}^{N} c(B_i, R_i) \leq \sum_{i=1}^{N} c(B_i, R_{i+1})$$

(where $R_{N+1} = R_1$). That is, something as in (1) does not occur.

Remark 6.2. In our discussion of bakeries and restaurants, we temporarily assign some label B_1 to some bakery. However, these labels are not considered as fixed, unless otherwise stated. That is, with respect to one labeling, Bob's Bakery is labeled as B_2, but with respect to another labeling, Bob's Bakery is labeled as B_5. So, in the above definition, I am allowed to change this labeling as I please. Hopefully this issue will not cause confusion.

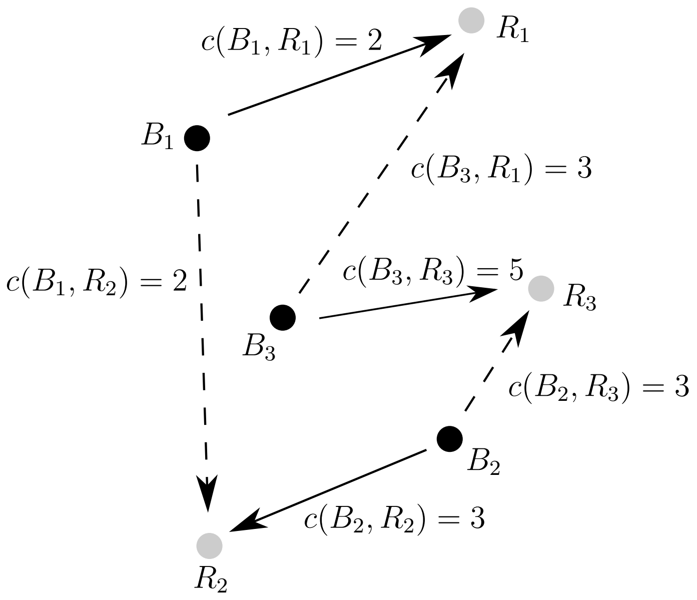

Figure 1. Solid lines denote the initial transference plan. For this plan, note that $c(B_1, R_1) + c(B_2, R_2) + c(B_3, R_3) = 10$. Dashed lines denote the adjusted transference plan. For this plan, note that $c(B_1, R_2)+c(B_2, R_3)+c(B_3, R_1) = 8$. Since the cost is less, the adjusted plan is better.

Let N be a positive integer, and suppose I have an optimal transportation plan. If my transportation plan is optimal, then whenever I transport loaves from B_1 to R_1, from B_2 to R_2, from ..., and from B_N to R_N, the set $(B_1, R_1), \ldots, (B_N, R_N)$ must be cyclically monotone. This assertion follows from our discussion above, as depicted in Figure 1. We now ask: is the converse of this assertion true? That is, can the cyclical monotonicity condition *guarantee* that my transportation plan is optimal? If so this would be nice.

To construct an optimal plan, I would then do the following. Look at some set $(B_1, R_1), \ldots,$ (B_N, R_N) of transported loaves such that B_1 ships a loaf to R_1, B_2 ships a loaf to R_2, ..., and B_N ships a loaf to R_N. If $\sum_{i=1}^{N} c(B_i, R_i) > \sum_{i=1}^{N} c(B_i, R_{i+1})$, then as discussed above, I can re-route the loaves to decrease my total cost. If $\sum_{i=1}^{N} c(B_i, R_i) \leq \sum_{i=1}^{N} c(B_i, R_{i+1})$, leave the loaves alone. Now look at a different set of transported loaves. Continue this checking and re-routing process. At each step, the cost is not increased, and there are only finitely many such sets to check. So, at the end of the process, my transportation plan would necessarily be optimal.

However, in the above paragraph, we have assumed that the converse of the above assertion holds. We do not yet know that this is true. Yet, it turns out to be true, as we will now demonstrate. Before we begin, we need to make precise our notion of an optimal plan. (We have been intentionally vague, for pedagogical reasons.) In the following discussion, the notation $\sum_{i,j=1}^{n}$ is shorthand for $\sum_{i=1}^{n} \sum_{j=1}^{n}$.

Definition 6.3. Fix some labeling of the bakeries and restaurants, so that $\mathcal{B} = (B_1, \ldots, B_n)$ and $\mathcal{R} = (R_1, \ldots, R_n)$. Let b_i be the number of loaves produced by bakery B_i ($i \in \{1, \ldots, n\}$), and let r_j be the number of loaves consumed by restaurant R_j ($j \in \{1, \ldots, n\}$). Assume that the total loaf production is equal to the total loaf consumption ($\sum_{i=1}^{n} b_i = \sum_{j=1}^{n} r_j$).

We define a **transportation plan** as a doubly indexed set of numbers $\{a_{i,j}\}_{i,j=1}^{n}$ such that $0 \leq a_{i,j} \leq 1$ for all $i,j \in \{1,\ldots,n\}$, $\sum_{i,j=1}^{n} a_{i,j} = 1$, $\sum_{j=1}^{n} a_{i,j} = b_i/(\sum_{k=1}^{n} b_k)$ for all $i \in \{1,\ldots,n\}$, and $\sum_{i=1}^{n} a_{i,j} = r_j/(\sum_{k=1}^{n} r_k)$ for all $j \in \{1,\ldots,n\}$.

Here $a_{i,j}$ represents the fraction of the total number of loaves $(\sum_{k=1}^{n} b_k)$ that I send from some bakery B_i to some restaurant R_j. For example, $b_i/(\sum_{k=1}^{n} b_k)$ is the fraction of loaves produced by bakery B_i, averaged with respect to the total production of all bakeries. Then the condition $\sum_{j=1}^{n} a_{i,j} = b_i/(\sum_{k=1}^{n} b_k)$ says that this fraction is equal to $\sum_{j=1}^{n} a_{i,j}$. Note that we do not require an integer number of loaves to be sent from a bakery to a restaurant. In particular, our above proof that an optimal transportation plan exists no longer applies. Recall that a doubly indexed set of numbers is called a **matrix**. From our definition, it is not obvious that a transportation plan exists at all. To see that at least one transportation plan exists, consider $a_{i,j} = b_i r_j/(\sum_{k=1}^{n} r_k)^2$.

Exercise 6.4. Check that $a_{i,j} = b_i r_j/(\sum_{k=1}^{n} r_k)^2$ is a transportation plan.

We now need to make a definition that applies to the transportation plan that says: whenever we transport some loaves, these loaves travel on a cyclically monotone set.

Definition 6.5. Fix some labeling of the bakeries and restaurants, so that $\mathcal{B} = (B_1,\ldots,B_n)$ and $\mathcal{R} = (R_1,\ldots,R_n)$. We say that a transportation plan $\{a_{i,j}\}_{i,j=1}^{n}$ is **cyclically monotone** if the following holds. Let $N \geq 1$ be any integer, and let $i(1),\ldots,i(N),j(1),\ldots,j(N) \in \{1,\ldots,n\}$. Let $(B_{i(1)}, R_{j(1))}),\ldots,(B_{i(N)}, R_{j(N)})$ be any set with $B_i \in \mathcal{B}, R_i \in \mathcal{R}$, $1 \leq i \leq N$. Assume that $a_{i(1),j(1)},\ldots,a_{i(N),j(N)} > 0$. Then the set $(B_{i(1)}, R_{j(1)}),\ldots,(B_{i(N)}, R_{j(N)})$ is cyclically monotone.

Let $A = \{(B_i, R_j) \in \mathcal{B}\times\mathcal{R} \colon a_{i,j} > 0\}$. A is called the **support** of the transportation plan, since it encompasses all paths of all transported loaves.

Definition 6.6. Fix some labeling of the bakeries and restaurants, so that $\mathcal{B} = (B_1,\ldots,B_n)$ and $\mathcal{R} = (R_1,\ldots,R_n)$. Let b_i be the number of loaves produced by bakery B_i ($i \in \{1,\ldots,n\}$), and let r_j be the number of loaves consumed by restaurant R_j ($j \in \{1,\ldots,n\}$). Suppose we have a transportation plan $\{a_{i,j}\}_{i,j=1}^{n}$. Then this transportation plan is called an **optimal transportation plan** if the following equation holds

$$\sum_{i,j=1}^{n} a_{i,j}\, c(B_i, R_j) = \min_{\substack{\{a'_{i,j}\}_{i,j=1}^{n} \colon \sum_{i,j=1}^{n} a'_{i,j}=1, \\ \sum_{j=1}^{n} a'_{i,j}=b_i/(\sum_{k=1}^{n} b_k), \forall i\in\{1,\ldots,n\}, \\ \sum_{i=1}^{n} a'_{i,j}=r_j/(\sum_{k=1}^{n} r_k), \forall j\in\{1,\ldots,n\}}} \sum_{i,j=1}^{n} a'_{i,j}\, c(B_i, R_j)$$

That is, this transportation plan minimizes the cost among all transportation plans.

Remark 6.7. In the rest of the discussion, the labeling of bakeries and restaurants is not fixed.

Theorem 6.8. ***(Optimal Transportation)*** *Suppose we have a cyclically monotone transportation plan $\{a_{i,j}\}_{i,j=1}^{n}$. Then $\{a_{i,j}\}_{i,j=1}^{n}$ is an optimal transportation plan.*

Proof. The strategy of the proof follows our discussion above. We define some function f that checks the condition of cyclical monotonicity. This function f will satisfy a special property, labeled $(**)$ below. Then a trick, known as duality, will deduce optimality of $\{a_{i,j}\}_{i,j=1}^{n}$ from $(**)$.

Let $A \subseteq \mathcal{B} \times \mathcal{R}$ be the support of the transportation plan $\{a_{i,j}\}_{i,j=1}^n$. For $B \in \mathcal{B}$, define a number $f(B) \in \mathbb{R}$ by the formula

$$\begin{aligned} f(B) = \max_{m \in \mathbb{N}} \max \{&[c(B_0, R_0) - c(B_1, R_0)] + [c(B_1, R_1) - c(B_2, R_1)] \\ &+ \cdots + [c(B_m, R_m) - c(B, R_m)] \colon (B_1, R_1), \ldots, (B_m, R_m) \in A\} \qquad (*) \end{aligned}$$

Let $(B_0, R_0) \in A$. If we take $(B_1, R_1) = (B_0, R_0)$ and $m = 1$ in the definition of $f(B_0)$, we see that $f(B_0) \geq [c(B_0, R_0) - c(B_0, R_0)] + [c(B_0, R_0) - c(B_0, R_0)] = 0$, so $f(B_0) \geq 0$. However, by cyclical monotonicity, each term of the form $[c(B_0, R_0) - c(B_1, R_0)] + [c(B_1, R_1) - c(B_2, R_1)] + \cdots + [c(B_m, R_m) - c(B_0, R_m)]$ is nonpositive. So, $f(B_0) \leq 0$. We conclude that $f(B_0) = 0$.

Now, by relabeling R_m as R, we write

$$\begin{aligned} f(B) = \max_{R \in \mathcal{R}} \max_{m \in \mathbb{N}} \max_{(B_1, R_1), \ldots, (B_{m-1}, R_{m-1}), B_m} &\{[c(B_0, R_0) - c(B_1, R_0)] + [c(B_1, R_1) - c(B_2, R_1)] \\ &+ \cdots + [c(B_{m-1}, R_{m-1}) - c(B_m, R_{m-1})] + [c(B_m, R) - c(B, R)] \\ &: (B_1, R_1), \ldots, (B_m, R) \in A\} \end{aligned}$$

For $R \in \mathcal{R}$ define a real number $g(R)$ by the formula

$$\begin{aligned} g(R) = \max \{&[c(B_0, R_0) - c(B_1, R_0)] + [c(B_1, R_1) - c(B_2, R_1)] \\ &+ \cdots + [c(B_{m-1}, R_{m-1}) - c(B_m, R_{m-1})] + c(B_m, R) \\ &: m \in \mathbb{N}, (B_1, R_1), \ldots, (B_{m-1}, R_{m-1}), (B_m, R) \in A\} \end{aligned}$$

(If the set $(B_1, R_1), \ldots, (B_{m-1}, R_{m-1}), (B_m, R) \in A$ is empty, define $g(R) = -\infty$.)

By the definitions of $f(B)$ and $g(R)$, we have

$$f(B) = \max_{R \in \mathcal{R}} (g(R) - c(B, R))$$

Let $(\overline{B}, \overline{R}) \in A$. In the original definition of f (equation $(*)$), choose $B_m = \overline{B}$ and $R_m = \overline{R}$. Note that $(B_1, R_1), \ldots, (B_m, R_m) \in A$ implies $(B_1, R_1), \ldots, (B_{m-1}, R_{m-1}) \in A$, so the definition of f implies that

$$\begin{aligned} f(B) &\geq \max_{m \in \mathbb{N}} \Bigg\{ \Bigg(\max_{(B_1, R_1), \ldots, (B_{m-1}, R_{m-1}) \in A} [c(B_0, R_0) - c(B_1, R_0)] \\ &\quad + \cdots + [c(B_{m-2}, R_{m-2}) - c(B_{m-1}, R_{m-2})] + [c(B_{m-1}, R_{m-1}) - c(\overline{B}, R_{m-1})] \Bigg) \\ &\qquad + [c(\overline{B}, \overline{R}) - c(B, \overline{R})] \Bigg\} \\ &= \max_{m \in \mathbb{N}} \Bigg\{ \Bigg(\max_{(B_1, R_1), \ldots, (B_m, R_m) \in A} [c(B_0, R_0) - c(B_1, R_0)] \\ &\quad + \cdots + [c(B_{m-1}, R_{m-1}) - c(B_m, R_{m-1})] + [c(B_m, R_m) - c(\overline{B}, R_m)] \Bigg) \Bigg\} \\ &\qquad + [c(\overline{B}, \overline{R}) - c(B, \overline{R})] \end{aligned}$$

Using the definition of f (equation $(*)$), the previous inequality says that

$$f(B) \geq f(\overline{B}) + c(\overline{B}, \overline{R}) - c(B, \overline{R})$$

That is, $f(B)+c(B,\overline{R}) \geq f(\overline{B})+c(\overline{B},\overline{R})$. Since this inequality holds for all $B \in \mathcal{B}$ and the right side does not depend on B, we can take the minimum of both sides over $B \in \mathcal{B}$ to get

$$\min_{B\in\mathcal{B}}(f(B)+c(B,\overline{R})) \geq f(\overline{B})+c(\overline{B},\overline{R})$$

However, taking $B = \overline{B}$ in the minimum of the left side, we see that $\min_{B\in\mathcal{B}}(f(B) + c(B,\overline{R})) \leq f(\overline{B})+c(\overline{B},\overline{R})$. Combining the two inequalities, we conclude

$$\min_{B\in\mathcal{B}}(f(B)+c(B,\overline{R})) = f(\overline{B})+c(\overline{B},\overline{R}) \qquad (**)$$

For $R \in \mathcal{R}$, define $h(R) = \min_{B\in\mathcal{B}}(f(B)+c(B,R))$. Then $(**)$ says that

$$h(\overline{R}) - f(\overline{B}) = c(\overline{B},\overline{R}), \qquad \forall\,(\overline{B},\overline{R}) \in A \qquad (\dagger)$$

Using that $a_{i,j}$ is a transference plan, and using the notation of Definition 6.3,

$$\begin{aligned}\frac{\sum_{j=1}^n r_j h(R_j)}{\sum_{k=1}^n r_k} - \frac{\sum_{i=1}^n b_i f(B_i)}{\sum_{k=1}^n r_k} &= \sum_{i,j=1}^n a_{i,j}h(R_j) - \sum_{i,j=1}^n a_{i,j} f(B_i) \\ &= \sum_{i,j=1}^n a_{i,j}[h(R_j)-f(B_i)] = \sum_{i,j=1}^n a_{i,j}c(B_i,R_j) \qquad \text{, by } (\dagger)\end{aligned}$$

Note our crucial use of the definition of A in the last equality.

Now, let ϕ, ψ be arbitrary functions such that $\phi\colon \mathcal{R} \to \mathbb{R}$, $\psi\colon \mathcal{B} \to \mathbb{R}$, and $\forall B \in \mathcal{B}, \forall R \in \mathcal{R}$, $\phi(R) - \psi(B) \leq c(B,R)$. Let $a'_{i,j}$ be an arbitrary transportation plan. Using that $a'_{i,j}$ is a transportation plan, we get

$$\begin{aligned}\frac{\sum_{j=1}^n r_j \phi(R_j)}{\sum_{k=1}^n r_k} - \frac{\sum_{i=1}^n b_i \psi(B_i)}{\sum_{k=1}^n r_k} &= \sum_{i,j=1}^n a'_{i,j}\phi(R_j) - \sum_{i,j=1}^n a'_{i,j}\psi(B_i) \\ &= \sum_{i,j=1}^n a'_{i,j}[\phi(R_j)-\psi(B_i)] \\ &\leq \sum_{i,j=1}^n a'_{i,j}c(B_i,R_j) \qquad \text{, since } \phi(R)-\psi(B) \leq c(B,R),\ \forall B\in\mathcal{B}, \forall R \in \mathcal{R}\end{aligned}$$

Since this inequality holds for all such ϕ, ψ and $a'_{i,j}$, we can take the maximum over such ϕ, ψ of both sides, and then take the minimum over such $a'_{i,j}$ of both sides to see that

$$\begin{aligned}&\max_{\substack{\phi,\psi\colon \phi(R)-\psi(B)\leq c(B,R),\\ \forall B\in\mathcal{B}, \forall R\in\mathcal{R}}} \left\{\frac{\sum_{j=1}^n r_j \phi(R_j)}{\sum_{k=1}^n r_k} - \frac{\sum_{i=1}^n b_i \psi(B_i)}{\sum_{k=1}^n r_k}\right\} \\ &\leq \min_{\substack{a'_{i,j}\colon a'_{i,j} \text{ is a} \\ \text{transportation plan}}} \left\{\sum_{i,j=1}^n a'_{i,j}c(B_i,R_j)\right\} \qquad (\ddagger)\end{aligned}$$

Finally, since $a'_{i,j} = a_{i,j}$, $\phi = h$, and $\psi = f$ satisfy the inequality $(\ddagger)$ with *equality*, both sides of $(\ddagger)$ are actually equal, and therefore $a_{i,j}$ achieves the minimum of the right side of $(\ddagger)$, as desired. □

Courant Institute, New York University, New York NY 10012
E-mail address: heilman@cims.nyu.edu