

PUTNAM MATH COMPETITION

PROF. OMAR

This book is a course in Putnam problem solving. It is organized around methods rather than around the usual undergraduate course titles. Each section gives intuition, highlighted principles, unboxed worked examples, and an unboxed practice set. The exercises are original Putnam-style problems unless otherwise noted. Putnam-labeled exercises are paraphrased from past exams, with official wording available through the MAA Putnam Archive. The solution chapters expand the proofs and give detailed solutions to selected problems, including every Putnam-labeled exercise.

CONTENTS

How to Use This Book	4
1 General Strategies for Putnam Problem Solving	5
2 Proof Techniques That Recur Constantly	8
3 Algebraic Manipulation and Structural Insight	11
4 Inequalities I: Standard Tools	14
5 Inequalities II: Olympiad and Putnam Techniques	17
6 Divisibility, Congruences, and Modular Arithmetic	21
7 Diophantine Equations and Integer Methods	24
8 Arithmetic Functions and Prime Factorization Techniques	27
9 Counting Methods and Double Counting	30
10 Pigeonhole Principle, Extremal Principles, and Probabilistic Thinking	33
11 Recurrences, Generating Functions, and Discrete Processes	36
12 Graph Theory for the Putnam	39
13 Euclidean Geometry and Transformational Methods	42
14 Coordinate Geometry, Vectors, and Analytic Methods	45
15 Linear Algebra Essentials for Problem Solving	48
16 Sequences, Series, and Asymptotic Thinking	51
17 Real Analysis Methods in Problem Solving	54
18 Calculus and Functional Methods	57

19	Polynomials, Complex Numbers, and Algebraic Methods	60
20	Functional Equations, Invariants, and Strategy Synthesis	63
A	Classical Theorems Every Putnam Student Should Know	66
B	Common Mistakes in Putnam Solutions	69
C	Timed Practice and Exam Strategy	72
D	A Roadmap Through Past Putnam Problems	76
E	Detailed Solutions and Proof Expansions	79
	Detailed solutions: Section 1	79
	Detailed solutions: Section 2	79
	Detailed solutions: Section 3	80
	Detailed solutions: Section 4	80
	Detailed solutions: Section 5	81
	Detailed solutions: Section 6	82
	Detailed solutions: Section 7	82
	Detailed solutions: Section 8	83
	Detailed solutions: Section 9	83
	Detailed solutions: Section 10	84
	Detailed solutions: Section 11	84
	Detailed solutions: Section 12	85
	Detailed solutions: Section 13	85
	Detailed solutions: Section 14	86
	Detailed solutions: Section 15	87
	Detailed solutions: Section 16	87
	Detailed solutions: Section 17	88
	Detailed solutions: Section 18	88
	Detailed solutions: Section 19	89
	Detailed solutions: Section 20	90
	Detailed solutions: Appendix exercises	90
F	Additional Worked Exercise Solutions	92
	Additional solutions: Section 1	92
	Additional solutions: Section 2	92
	Additional solutions: Section 3	93
	Additional solutions: Section 4	93

Additional solutions: Section 5	93
Additional solutions: Section 6	94
Additional solutions: Section 7	94
Additional solutions: Section 8	95
Additional solutions: Section 9	95
Additional solutions: Section 10	96
Additional solutions: Section 11	96
Additional solutions: Section 12	97
Additional solutions: Section 13	97
Additional solutions: Section 14	97
Additional solutions: Section 15	98
Additional solutions: Section 16	98
Additional solutions: Section 17	99
Additional solutions: Section 18	99
Additional solutions: Section 19	99
Additional solutions: Section 20	100
G Further Proof Models	101
Expanded solutions to revised exercises	101
Topic-based proof models	106

How to Use This Book

This book is designed for active work. Read each section with paper beside you. When an example begins, pause before reading the solution and write the first move you would try. After reading the solution, summarize the method in one sentence. A useful sentence has a verb: choose a minimal counterexample, count incidences, normalize the variables, pass to residues, or encode the process as a graph.

The 2026 competition format

Beginning with the 2026 competition, the six hours of competition time are divided into *four ninety-minute sessions* rather than two three-hour sessions. The sessions have more frequent designated breaks. During a session, a student who exits the exam room may not return to continue working in that session. The MAA is also making start times synchronous across time zones. Because the detailed timetable is supplied through the official FAQ and the forthcoming supervisor instructions, verify the reporting time and break rules with your local supervisor rather than relying on the schedule used in an earlier year.^a

^aMathematical Association of America, *MAA Putnam*, 2026 update, <https://maa.org/putnam/>.

This change affects pacing but not the central mathematical habits in the book. Training should now emphasize rapid orientation inside a ninety-minute packet, complete proof writing before each session ends, deliberate use of the designated breaks, and practical preparation before entering the room.

A four-pass reading method

Pass 1: orientation. Read the intuition paragraphs and theorem statements. Do not worry about mastering every detail.

Pass 2: examples. Work through the examples line by line, filling in omitted algebra.

Pass 3: exercises. Attempt the first four exercises for fluency and the last two for depth.

Pass 4: compression. Write a half-page summary of the section with triggers, favorite examples, and common traps.

The exercises are arranged to increase in difficulty. The last problems in a section are meant to feel like contest problems: they may require a choice of method rather than direct application of the preceding theorem. Do not read difficulty as a judgment. A problem is difficult when its first move is hidden, not when it uses a long list of facts.

What to write in a solution journal

For each problem, record the problem source, the first useful observation, the key tool, the equality or extremal case, and the final proof spine. The proof spine is the shortest sequence of ideas that makes the solution work. Reviewing proof spines is one of the fastest ways to prepare for mixed Putnam sets.

§1 General Strategies for Putnam Problem Solving

A Putnam problem rarely asks for a routine computation. It usually gives a plain object and hides a structure inside it. The central skill is not knowing a theorem first. The central skill is deciding what the object wants to become. A set might want to become a graph, a recurrence might want a generating function, an inequality might want a convexity statement, and a divisibility question might want a congruence class.

A useful way to begin is to ask for the **small cases**, the **extreme cases**, and the **invariants**. Small cases suggest the answer. Extreme cases show where equality or contradiction is likely to live. Invariants tell which quantities cannot change during a process. The same habit appears throughout this book, so our recurring examples will be finite sets, polynomials with prescribed roots, and directed graphs encoding choices.

The Putnam first minute

Before writing a proof, spend a minute on the following questions. What happens for the smallest values? What quantity is forced to be integer, positive, monotone, or bounded? What changes if the problem is reversed? What would an equality case look like? What object could encode the information more cleanly?

Example 1.1. A hidden invariant in a grid walk

A token starts at $(0, 0)$ in the integer lattice. A move replaces (x, y) by one of $(x + 1, y + 2)$, $(x + 2, y + 1)$, or $(x - 1, y - 1)$. Can the token ever reach $(1, 1)$?

Solution. Looking only at coordinates is distracting because the move set is not symmetric. The sum $x + y$ changes by 3, 3, or -2 . Modulo 5, these changes are 3, 3, 3, so every move changes $x + y$ by 3 modulo 5. After n moves we have $x + y \equiv 3n \pmod{5}$. This alone does not rule out $(1, 1)$ because $2 \equiv 3 \cdot 4 \pmod{5}$. A second observation is better: the difference $x - y$ changes by -1 , 1, or 0. The first two moves increase $x + y$ by 3, while the third decreases it by 2 and leaves $x - y$ fixed. Suppose the endpoint is $(1, 1)$. If the three moves are used a, b, c times, then

$$a + 2b - c = 1, \quad 2a + b - c = 1.$$

Subtracting gives $b = a$. Substituting gives $3a - c = 1$, so $c = 3a - 1$. The number of moves is $a + b + c = 5a - 1$. But the sum $x + y = 2$ also equals $3a + 3b - 2c = 2$, which is consistent. The obstruction must be nonnegativity: $c = 3a - 1$ requires $a \geq 1$, and for $a = 1$ we can use one of each positive move and two negative moves. Indeed a route exists, for instance $(0, 0) \rightarrow (1, 2) \rightarrow (3, 3) \rightarrow (2, 2) \rightarrow (1, 1)$. The lesson is as important as the answer: a suspected invariant must be tested against variables that count the moves.

The example is deliberately imperfect. Many invariants fail to decide the problem. On the Putnam, a failed invariant is still useful because it tells you what data the answer depends on. A strong solution often comes after two weak attempts.

Example 1.2. Extremal choice in a finite set

Let S be a nonempty finite set of positive integers such that whenever $a, b \in S$ and $a > b$, the integer $a - b$ also lies in S . Prove that S consists of all positive multiples of its least element up to its largest element.

Solution. Let m be the least element of S . First observe that every element of S is divisible by m . If not, choose $a \in S$ with least positive remainder modulo m . Write $a = qm + r$ with $0 < r < m$. Because $a, m \in S$, repeated subtraction gives $a - m, a - 2m, \dots, r \in S$, contradicting minimality of m . Now let M be the largest element. If $km \leq M$, write $M = tm$. Starting from M and subtracting m repeatedly shows $(t - 1)m, (t - 2)m, \dots, m \in S$. Thus every positive multiple of m up to M is present.

The proof used two standard moves: choose the least counterexample and reduce it, then use closure to fill a whole interval. This pattern will return in number theory, graph theory, and analysis.

Three universal moves

- **Normalize.** Divide by a scale, translate the origin, sort the variables, or set the smallest object equal to a simple symbol.
- **Encode.** Replace a relation by a graph, a sequence by a generating function, a congruence by a residue class, or a geometry problem by coordinates.
- **Reverse.** Ask what a desired conclusion would force. Many Putnam problems are easier from the end than from the beginning.

The experiment to proof pipeline

The usual path is not guess, then proof. It is experiment, pattern, obstruction, proof. Make a table of small cases, but do not stop at the table. After the table, write one sentence explaining what feature of the small cases survives. That sentence becomes the proof target. For example, if a process seems to preserve parity, the proof target is not the final answer; it is the sentence “this parity is invariant under every move.”

Example 1.3. A chessboard invariant

A standard chessboard has two opposite corner squares removed. Can it be tiled by dominoes, each covering two adjacent squares?

Solution. Color the board in the usual black-white pattern. Opposite corners have the same color, so after removing them the board has 30 squares of one color and 32 of the other. Every domino covers one black square and one white square. Therefore any domino tiling would cover equal numbers of black and white squares, impossible. The invariant is not the number of squares but the difference between black and white squares.

Example 1.4. Working backward

A positive integer is called reachable if it can be obtained from 1 by repeatedly applying either $n \mapsto 2n$ or $n \mapsto 3n + 1$. Prove that every reachable integer is odd only if it is 1 or has a predecessor of the form $(m - 1)/3$.

Solution. Instead of trying to generate all reachable numbers, look at the last move. If a reachable integer m is odd and $m \neq 1$, it cannot have been obtained by the move $n \mapsto 2n$, since that output is even. Hence the last move must have been $n \mapsto 3n + 1$, so $m = 3n + 1$ and $n = (m - 1)/3$ is a reachable integer. This tiny example illustrates a powerful habit: when the forward process branches, the last step may be forced by a parity, size, or divisibility feature of the target.

A running notebook format

For each serious problem, keep four lines of notes: the small cases, the guessed invariant, the failed idea, and the final trigger. Over time the failed ideas become as valuable as the solved problems. They teach which signals are weak and which signals are reliable.

Problem patterns for this section

General strategy problems often hide their method in the verbs. “Can this be reached” suggests invariants or reverse moves. “Show there exists” suggests pigeonhole, averaging, compactness, or construction. “Find

all” demands two halves: prove every answer has a necessary form, then construct or verify every object of that form. In practice, write the two halves as separate paragraphs. This prevents a common error where a solution proves only that a proposed answer is plausible.

Proof-writing details for this section

In a strategy proof, the most important details are the ones that justify the choice of viewpoint. Do not merely announce an invariant or an extremal element. State exactly how each legal move changes the proposed quantity, and state exactly why the initial and terminal values are incompatible or compatible. When an extremal choice is made, identify the ordered set in which the extremum lives and explain why the operation produces a strictly smaller or larger object inside the same set.

Exercises for Section 1

- (1) A sequence starts with $a_0 = 1$ and each move replaces the current value a by either $2a + 1$ or $3a + 1$. Prove that every reachable integer other than 1 is either odd or congruent to 1 modulo 3. Explain why the last move makes this condition unavoidable.
- (2) Let S be a finite set of real numbers. Suppose that for every $x, y \in S$, the average $(x + y)/2$ lies in S . Prove that S has one element.
- (3) On an $m \times n$ board, a move flips every square in one chosen row and then every square in one chosen column; the intersection square is therefore flipped twice and returns to its original color. Prove that when $m + n$ is even, the parity of the number of black squares is invariant.
- (4) Let $a_1, \dots, a_n$ be distinct positive integers. Show that some nonempty subset has sum divisible by n .
- (5) A function $f : \mathbb{N} \rightarrow \mathbb{N}$ satisfies $f(f(n)) = n + 1$ for all n . Prove that no such function exists.
- (6) Let $n \geq 2$ and let $x_1, \dots, x_n$ be real numbers with sum 0. Prove that there is a nonempty proper subset whose average is at least 0.
- (7) (**Putnam 2001 B1**) Let n be even. Fill an $n \times n$ grid row by row with $1, 2, \dots, n^2$. Color exactly half of the squares in each row and each column red, and the rest black. Prove that the sum of the red entries equals the sum of the black entries.

§2 Proof Techniques That Recur Constantly

Proof techniques are not decorations. They are ways of controlling uncertainty. Induction controls infinitely many cases by forcing a local transition. Contradiction controls a statement by imagining the first place where it fails. The contrapositive controls implication by moving the hard conclusion into the hypothesis. The Putnam rewards proofs that choose the form which makes the obstruction visible.

The same theorem can have many proof forms. If a statement says every object has a property, induction or contradiction are natural. If it says no object has a property, parity, bounding, and infinite descent are natural. If it says at least one object exists, averaging, pigeonhole, compactness, or construction are natural.

Theorem 2.1 Strong induction as compression

Suppose $P(n)$ is a statement about positive integers. If $P(1)$ holds and, for every $n > 1$, the truth of all $P(k)$ with $1 \leq k < n$ implies $P(n)$, then $P(n)$ holds for all $n \geq 1$.

The point of strong induction is not that it is stronger logically. It lets you use the natural smaller object, even when the smaller object has unpredictable size. Factorization, graph deletion, recurrence splitting, and Euclidean algorithm arguments often need this freedom.

Example 2.1. Induction with a removable object

Every positive integer n can be written as a sum of distinct powers of 2.

Proof. The statement is clear for $n = 1$. For $n > 1$, let 2^k be the largest power of 2 not exceeding n . If $n = 2^k$ there is nothing to prove. Otherwise $0 < n - 2^k < 2^k \leq n$, so by the induction hypothesis $n - 2^k$ is a sum of distinct powers of 2. None of those powers can be 2^k because their sum is smaller than 2^k . Adding 2^k gives the required representation.

This proof illustrates the typical induction plan: remove the largest clean piece and verify that the remainder cannot interfere with what was removed.

Lemma 2.1 Minimal counterexample principle

If a statement about positive integers is false, then there is a least positive integer for which it is false.

Proof. This is exactly the well-ordering property of the positive integers. A nonempty subset of $\mathbb{N}$ has a least element, so the set of counterexamples, if nonempty, has a least element.

Example 2.2. Contradiction by infinite descent

Prove that $\sqrt{2}$ is irrational.

Proof. Suppose $\sqrt{2} = a/b$ with positive integers a, b and with b as small as possible. Then $a^2 = 2b^2$, so a is even. Write $a = 2c$. Then $4c^2 = 2b^2$, so $b^2 = 2c^2$ and b is even. Thus $a/b = c/(b/2)$ gives the same rational number with smaller positive denominator, contradicting the choice of b .

Example 2.3. A bijective proof of a binomial identity

Prove that

$$\sum_{k=0}^n \binom{n}{k}^2 = \binom{2n}{n}.$$

Solution. Choose n people from a group of $2n$ people divided into two teams of n each. The right side counts the choice directly. If exactly k people are chosen from the first team, then $n - k$ are chosen from the second team, so the number of such selections is $\binom{n}{k}\binom{n}{n-k} = \binom{n}{k}^2$. Summing over k counts the same selections by cases.

A bijection or double count is often cleaner than algebra because it explains why the identity should be true. Whenever a sum contains binomial coefficients, ask what it counts.

Choosing the proof form

Induction	Look for a way to delete, contract, factor, or reduce dimension.
Contradiction	Assume the obstruction exists and choose the smallest, largest, nearest, or most extreme obstruction.
Contrapositive	Use when the conclusion is negative or when failure of the conclusion has concrete structure.
Construction	Build the object in stages and preserve an invariant.
Bijection	Use when two expressions count the same hidden set.

Example 2.4. Contrapositive as a cleaner implication

Prove that if n is an integer and n^2 is odd, then n is odd.

Proof. The contrapositive is easier: if n is not odd, then n is even. Write $n = 2k$. Then $n^2 = 4k^2 = 2(2k^2)$ is even, so n^2 is not odd. Therefore the original implication holds.

The contrapositive works well whenever the conclusion has a simpler negation than the original hypothesis. “Not relatively prime” means a common prime divisor exists. “Not injective” means two distinct inputs have the same output. “Not connected” means a partition into two noncommunicating pieces exists.

Example 2.5. Construction with a maintained condition

Construct, for each n , a set of n distinct positive integers such that no one divides another.

Solution. Choose the integers

$$n + 1, n + 2, \dots, 2n.$$

If a and b are two distinct numbers in this interval and $a < b$, then $2a \geq 2(n + 1) > 2n \geq b$. Thus b cannot be a multiple of a other than a itself. Since the numbers are distinct, no one divides another. The construction works because the interval is short on a multiplicative scale.

Proof repair questions

After writing a proof, ask whether every variable has been defined, whether every division has a nonzero denominator, whether every chosen minimum or maximum exists, and whether the proof handles equality cases. These questions are mechanical, but they prevent most lost points.

Example 2.6. Induction on structure instead of number

Prove that every expression formed from real numbers using addition and multiplication has the same value after any legal insertion of parentheses that preserves the order of the numbers and operations.

Proof. The natural induction is on the number of operations in the expression. If there are no operations, there is nothing to prove. Otherwise the outermost operation splits the expression into a left subexpression and a

right subexpression. By induction, the values of those two subexpressions are independent of their internal parenthesizing. Applying the final operation then gives a value independent of the full parenthesizing. This is structural induction: we reduce an object by splitting it into smaller objects of the same type.

Problem patterns for this section

When a problem contains a positive integer parameter, induction is a candidate but not a guarantee. Ask what the smaller object is. If you cannot name it, induction may be the wrong tool. When a problem says an object is impossible, contradiction is natural, but the contradiction should use a minimal counterexample, a parity obstruction, or a violated bound. A proof technique is useful only when it gives you a concrete object to manipulate.

Proof-writing details for this section

A proof form should be visible to the reader before the algebra begins. In induction, name the induction statement and the smaller object being used. In contradiction, name the minimal counterexample or the impossible obstruction. In a bijective proof, describe both directions of the map and explain why they undo each other. These sentences prevent the proof from becoming a list of correct but disconnected computations.

Exercises for Section 2

- (1) Prove by induction that $1 + 2 + \cdots + n = n(n+1)/2$.
- (2) Prove that every integer $n \geq 2$ is a product of primes using strong induction.
- (3) Prove that there are infinitely many primes using contradiction.
- (4) Prove that if n^2 is divisible by 3, then n is divisible by 3, and use the result to prove that $\sqrt{3}$ is irrational.
- (5) Give a bijective proof of $\sum_{k=0}^n \binom{n}{k} = 2^n$.
- (6) Let $a_1, \dots, a_n$ be integers. Prove that there exist two partial sums $a_1 + \cdots + a_i$ and $a_1 + \cdots + a_j$ with difference divisible by n , or one partial sum divisible by n .
- (7) (**Putnam 1999 A5**) Prove that there is a constant C such that every polynomial $p(x)$ of degree 1999 satisfies $|p(0)| \leq C \int_{-1}^1 |p(x)| dx$.

§3 Algebraic Manipulation and Structural Insight

Algebra in Putnam problems is not just expansion. The goal is to expose a hidden factor, symmetry, or monotone quantity. A good manipulation makes the problem shorter. A bad manipulation makes it longer without changing its shape. When you see symmetric expressions, try elementary symmetric sums. When you see fractions, try clearing denominators only after deciding what should factor. When you see a parameter, try making the expression homogeneous.

Definition 3.1 Homogeneity

An expression in variables $x_1, \dots, x_n$ is homogeneous of degree d if replacing each x_i by tx_i multiplies the expression by t^d .

Homogeneity matters because it allows normalization. If a statement is homogeneous, we may often impose $x + y + z = 1$, $xyz = 1$, or $\sum x_i^2 = 1$ without losing generality. This reduces one degree of freedom and often reveals convexity or a standard inequality.

Example 3.1. A factor that wants to be seen

Factor

$$x^3 + y^3 + z^3 - 3xyz.$$

Use the factorization to prove that if $x + y + z = 0$, then $x^3 + y^3 + z^3 = 3xyz$.

Solution. The expression is symmetric and vanishes when $x = y = z$, so a factor involving $x + y + z$ is plausible. The standard factorization is

$$x^3 + y^3 + z^3 - 3xyz = (x + y + z)(x^2 + y^2 + z^2 - xy - yz - zx).$$

One verifies it by expansion. If $x + y + z = 0$, the first factor is zero, so the identity gives $x^3 + y^3 + z^3 - 3xyz = 0$.

The second factor can also be written as

$$\frac{1}{2}((x - y)^2 + (y - z)^2 + (z - x)^2),$$

which reveals nonnegativity. A factorization with a sum of squares is especially valuable in inequality and equality-case problems.

Theorem 3.1 Polynomial identity principle

If two polynomials of degree at most d agree at more than d distinct points, then they are identical.

Proof. The difference of the two polynomials has degree at most d and has more than d roots. A nonzero polynomial of degree at most d cannot have more than d roots, so the difference is the zero polynomial.

Example 3.2. Interpolation without finding the polynomial

Let $P(x)$ be a polynomial of degree at most n such that $P(k) = 0$ for $k = 0, 1, \dots, n - 1$ and $P(n) = 1$. Find $P(n + 1)$.

Solution. Since P has roots $0, 1, \dots, n - 1$, it has the form

$$P(x) = cx(x - 1) \cdots (x - n + 1).$$

Using $P(n) = 1$ gives $c = 1/n!$. Therefore

$$P(n+1) = \frac{(n+1)n(n-1)\cdots 2}{n!} = n+1.$$

The computation succeeds because we let the roots determine the structure before expanding anything.

Lemma 3.1 Symmetric quadratic reduction

For real numbers x, y, z ,

$$x^2 + y^2 + z^2 - xy - yz - zx \geq 0,$$

with equality exactly when $x = y = z$.

Proof. Rewrite the left side as $\frac{1}{2}((x-y)^2 + (y-z)^2 + (z-x)^2)$. It is a sum of squares, so it is nonnegative. Equality holds only when all three squares vanish, which is exactly $x = y = z$.

Algebraic smell tests

- If an expression is symmetric, try elementary symmetric sums or equality cases.
- If it vanishes under a simple condition, look for the corresponding factor.
- If it is homogeneous, normalize one natural quantity.
- If it has many prescribed values, use roots or interpolation.
- If it has square roots, isolate only after checking for sign conditions.

Example 3.3. Completing the square

Find the minimum value of $x^2 - 6x + 13$ over all real x .

Solution. Completing the square gives

$$x^2 - 6x + 13 = (x - 3)^2 + 4.$$

The square term is nonnegative, so the minimum is 4, attained at $x = 3$. The method works because a quadratic is best understood by its vertex rather than by its expanded form.

Example 3.4. Telescoping by algebraic rewriting

Evaluate

$$\sum_{k=1}^n \frac{1}{k(k+1)(k+2)}.$$

Solution. The denominator suggests a difference of simpler reciprocal products:

$$\frac{1}{k(k+1)(k+2)} = \frac{1}{2} \left(\frac{1}{k(k+1)} - \frac{1}{(k+1)(k+2)} \right).$$

Thus the sum telescopes to

$$\frac{1}{2} \left(\frac{1}{1 \cdot 2} - \frac{1}{(n+1)(n+2)} \right) = \frac{1}{4} - \frac{1}{2(n+1)(n+2)}.$$

The manipulation is not random partial fractions. It looks for neighboring terms that cancel after summation.

When to expand and when not to expand

Do not expand products unless expansion reveals coefficients, roots, or cancellation. Keep factored forms for divisibility, sign, and roots. Keep squared forms for nonnegativity. Keep symmetric forms for substitutions such as $s = x + y + z$, $p = xy + yz + zx$, and $q = xyz$.

Example 3.5. A symmetric substitution

Suppose $x + y = 3$ and $xy = 2$. Compute $x^3 + y^3$ without finding x and y .

Solution. Use the identity

$$x^3 + y^3 = (x + y)^3 - 3xy(x + y).$$

Substitution gives $27 - 18 = 9$. This style is useful when the variables are roots of a polynomial and only symmetric information is available.

Problem patterns for this section

Algebra problems often become short after one normalization. If a polynomial has many roots, divide by the root factor. If an expression has equality when variables are equal, factor out differences such as $(x - y)^2$. If an expression has a denominator, decide whether clearing denominators preserves sign. If a problem gives values at integers, try finite differences or interpolation before expanding unknown coefficients.

Proof-writing details for this section

Algebraic proofs need verification at the points where a transformation can lose information. If denominators are cleared, record that they are nonzero. If a factorization is used, explain how it was found and why the remaining factors have the claimed sign or divisibility. For polynomials, state the degree bound and the number of roots before concluding that a polynomial is identically zero.

Exercises for Section 3

- (1) Factor $x^4 - y^4$ in two different useful ways.
- (2) Prove that $a^3 - b^3$ is divisible by $a - b$ for all integers a, b .
- (3) Let P be a polynomial of degree at most 3 with $P(0) = 0$, $P(1) = 1$, $P(2) = 4$, and $P(3) = 9$. Find $P(4)$.
- (4) Prove that $x^4 + y^4 \geq 2x^2y^2$ and determine equality.
- (5) Let a, b, c be real numbers with $a + b + c = 0$. Prove that $a^4 + b^4 + c^4 = 2(a^2b^2 + b^2c^2 + c^2a^2)$.
- (6) Let P be a monic polynomial of positive degree whose roots are all integers. Prove that if $P(k) = \pm 1$ for an integer k , then every root of P belongs to $\{k - 1, k + 1\}$. Conversely, determine exactly which multiplicities of the two roots give $P(k) = 1$ and which give $P(k) = -1$.
- (7) (**Putnam 1999 A2**) Let $p(x)$ be a real polynomial that is nonnegative for every real x . Prove that p is a finite sum of squares of real polynomials.

§4 Inequalities I: Standard Tools

Inequalities are often about choosing the right language. AM-GM speaks the language of products and sums. Cauchy-Schwarz speaks the language of squares and dot products. Jensen speaks the language of convexity. Rearrangement speaks the language of sorted order. On the Putnam, the hard part is rarely applying a theorem after it is visible. The hard part is transforming the expression until the theorem is visible.

Theorem 4.1 AM-GM inequality

For nonnegative real numbers $x_1, \dots, x_n$,

$$\frac{x_1 + \dots + x_n}{n} \geq (x_1 x_2 \dots x_n)^{1/n},$$

with equality if and only if $x_1 = \dots = x_n$.

Proof. A standard proof begins with the two-variable case, where $(\sqrt{x} - \sqrt{y})^2 \geq 0$ gives $(x + y)/2 \geq \sqrt{xy}$. Repeated averaging proves the result for powers of two, and a limiting or padding argument gives the general case. The equality condition persists through the averaging: equality is possible only when every averaged pair is equal, so all variables are equal.

Example 4.1. A product with fixed sum

Let $x, y, z > 0$ with $x + y + z = 3$. Prove that $xyz \leq 1$.

Solution. AM-GM gives

$$1 = \frac{x + y + z}{3} \geq (xyz)^{1/3}.$$

Cubing yields $xyz \leq 1$, with equality at $x = y = z = 1$. The equality case was visible from the condition $x + y + z = 3$, which is a useful guide: when a symmetric inequality has a fixed sum, equal variables are usually the first test.

Theorem 4.2 Cauchy-Schwarz inequality

For real numbers a_i, b_i ,

$$\left(\sum_{i=1}^n a_i b_i \right)^2 \leq \left(\sum_{i=1}^n a_i^2 \right) \left(\sum_{i=1}^n b_i^2 \right).$$

Equality holds exactly when the two vectors are proportional.

Proof. Consider the quadratic polynomial $\sum_i (a_i t - b_i)^2$ in t . It is nonnegative for every real t , so its discriminant is nonpositive. Expanding the discriminant gives exactly the stated inequality.

Example 4.2. Cauchy in Engel form

For positive real numbers $x_1, \dots, x_n$, prove that

$$\sum_{i=1}^n \frac{1}{x_i} \geq \frac{n^2}{x_1 + \dots + x_n}.$$

Solution. Write

$$\left(\sum_{i=1}^n 1 \right)^2 = \left(\sum_{i=1}^n \frac{1}{\sqrt{x_i}} \sqrt{x_i} \right)^2.$$

Cauchy gives

$$n^2 \leq \left(\sum_{i=1}^n \frac{1}{x_i} \right) \left(\sum_{i=1}^n x_i \right),$$

and the result follows. This form is often called Engel form or Titu Andreescu's lemma when written as $\sum a_i^2/x_i \geq (\sum a_i)^2/\sum x_i$.

Theorem 4.3 Jensen's inequality

If f is convex on an interval and $x_1, \dots, x_n$ lie in that interval, then

$$f\left(\frac{x_1 + \dots + x_n}{n}\right) \leq \frac{f(x_1) + \dots + f(x_n)}{n}.$$

The inequality reverses for concave functions.

Jensen is most useful when the problem has a fixed average and a sum of identical functions. For example, sums of logarithms point to products, sums of exponentials point to convexity, and sums of square roots point to concavity.

Theorem 4.4 Rearrangement inequality

If $a_1 \leq \dots \leq a_n$ and $b_1 \leq \dots \leq b_n$, then

$$\sum_{i=1}^n a_i b_{n+1-i} \leq \sum_{i=1}^n a_i b_{\sigma(i)} \leq \sum_{i=1}^n a_i b_i$$

for every permutation σ .

Proof. It is enough to show that swapping an inverted pair cannot decrease the sum. If $a_i \leq a_j$ but $b_{\sigma(i)} > b_{\sigma(j)}$, then replacing the two assigned b values changes the sum by

$$(a_i b_{\sigma(j)} + a_j b_{\sigma(i)}) - (a_i b_{\sigma(i)} + a_j b_{\sigma(j)}) = (a_j - a_i)(b_{\sigma(i)} - b_{\sigma(j)}) \geq 0.$$

Repeating swaps removes all inversions and reaches the similarly sorted pairing.

Example 4.3. Chebyshev's inequality for similarly sorted sequences

Let $a_1 \leq \dots \leq a_n$ and $b_1 \leq \dots \leq b_n$. Prove that

$$\frac{1}{n} \sum a_i b_i \geq \left(\frac{1}{n} \sum a_i \right) \left(\frac{1}{n} \sum b_i \right).$$

Solution. Expand the difference between the two sides after multiplying by n^2 :

$$n \sum_i a_i b_i - \sum_i \sum_j a_i b_j = \sum_{i < j} (a_j - a_i)(b_j - b_i) \geq 0.$$

The sorted order is the whole reason the inequality is true. Without it, the statement can fail.

Equality cases as a compass

When a standard inequality is used, record the equality condition immediately. In AM-GM, equality means the averaged quantities are equal. In Cauchy, equality means proportional vectors. In Jensen, equality usually means equal inputs or linear behavior of the function on the relevant interval. Equality conditions often tell you whether the chosen inequality is sharp enough.

Example 4.4. A Cauchy equality check

For positive x, y, z with $x + y + z = 1$, prove

$$\frac{x^2}{x+y} + \frac{y^2}{y+z} + \frac{z^2}{z+x} \geq \frac{1}{2}.$$

Solution. By Engel form of Cauchy,

$$\sum_{cyc} \frac{x^2}{x+y} \geq \frac{(x+y+z)^2}{(x+y) + (y+z) + (z+x)} = \frac{1}{2}.$$

Equality in Cauchy requires $x/(x+y) = y/(y+z) = z/(z+x)$ in the proportional form, and $x = y = z = 1/3$ indeed gives equality.

Problem patterns for this section

Standard inequality problems usually have a visible equality case. Check it first. If equality occurs when all variables are equal, AM-GM, Cauchy, Jensen, or smoothing may be appropriate. If equality occurs at the boundary, a convexity or extremal argument may be better. If the expression contains reciprocals, try Engel form. If it contains a product under a fixed sum, try AM-GM or logarithms.

Proof-writing details for this section

An inequality proof should keep its equality case visible. After applying AM-GM, Cauchy, Jensen, or rearrangement, write the equality condition immediately and compare it with the problem's constraints. If the equality condition does not match the expected extremal case, the inequality may still be useful, but it is probably too weak to finish the problem alone.

Exercises for Section 4

- (1) Prove that $x^2 + y^2 \geq 2xy$ for real x, y .
- (2) Let $x, y > 0$ and $xy = 1$. Prove that $x + y \geq 2$.
- (3) Prove that for positive a, b, c , $\frac{1}{a} + \frac{1}{b} + \frac{1}{c} \geq \frac{9}{a+b+c}$.
- (4) Let $x_1, \dots, x_n > 0$ with product 1. Prove that $x_1 + \dots + x_n \geq n$.
- (5) Let $x, y, z > 0$ with $x + y + z = 1$. Prove that $\frac{x}{y+z} + \frac{y}{z+x} + \frac{z}{x+y} \geq \frac{3}{2}$.
- (6) Let $x_1, \dots, x_n \in (0, \pi)$ with fixed sum. Find the maximum possible value of $\prod_i \sin x_i$.
- (7) (**Putnam 2003 A2**) Let $a_1, \dots, a_n$ and $b_1, \dots, b_n$ be nonnegative real numbers. Prove that $(a_1 \cdots a_n)^{1/n} + (b_1 \cdots b_n)^{1/n} \leq ((a_1 + b_1) \cdots (a_n + b_n))^{1/n}$.

§5 Inequalities II: Olympiad and Putnam Techniques

The second layer of inequalities is strategic. Instead of asking which named inequality applies, ask how the expression changes when variables are mixed, smoothed, or normalized. Many symmetric inequalities are solved by proving that the extreme case occurs when variables are equal or when some variables collide. The Putnam often uses this idea in a disguised way: an optimization problem becomes a statement about where a function has its maximum.

Definition 5.1 Smoothing

Smoothing is the method of replacing two variables by a more balanced pair with the same sum, product, or another conserved quantity, then showing that the expression moves in the desired direction.

Smoothing is powerful because it turns an n -variable problem into a two-variable check. If repeated smoothing improves the expression, the optimum is reached at equal variables or at the boundary.

Example 5.1. Smoothing a sum of squares

Let $x, y, z \geq 0$ with $x + y + z = 1$. Find the maximum and minimum of $x^2 + y^2 + z^2$.

Solution. The minimum occurs when the variables are as equal as possible. Indeed, replacing x, y by their average keeps the sum fixed and changes the contribution by

$$x^2 + y^2 - 2 \left(\frac{x+y}{2} \right)^2 = \frac{(x-y)^2}{2} \geq 0.$$

Repeated averaging lowers the sum of squares until $x = y = z = 1/3$, giving minimum $1/3$. For the maximum, convexity says a convex function on a simplex is maximized at a vertex. Here the maximum is 1, attained at $(1, 0, 0)$ and its permutations.

Lemma 5.1 Tangent line method

If f is convex, then every tangent line lies below the graph. If f is concave, every tangent line lies above the graph.

Proof. For a differentiable convex function, the slope of secant lines is increasing. Thus for any fixed a , the graph lies above $f(a) + f'(a)(x - a)$. The concave case follows by applying the convex case to $-f$.

Example 5.2. A tangent line bound

Prove that $\ln x \leq x - 1$ for all $x > 0$.

Solution. The function $\ln x$ is concave, so its tangent line at $x = 1$ lies above the graph. The tangent line is $y = x - 1$. Hence $\ln x \leq x - 1$, with equality only at $x = 1$. This one-line inequality is a Putnam staple because it converts products into sums.

Example 5.3. Normalization and equality cases

Let $a, b, c > 0$ and $abc = 1$. Prove that

$$\frac{a}{a^2 + a + 1} + \frac{b}{b^2 + b + 1} + \frac{c}{c^2 + c + 1} \leq 1.$$

Solution. Because $abc = 1$, write $a = x/y$, $b = y/z$, and $c = z/x$ for positive x, y, z . Then

$$\frac{a}{a^2 + a + 1} = \frac{x/y}{x^2/y^2 + x/y + 1} = \frac{xy}{x^2 + xy + y^2}.$$

The desired sum becomes

$$\sum_{cyc} \frac{xy}{x^2 + xy + y^2} \leq 1.$$

Since $x^2 + y^2 \geq 2xy$, we have $x^2 + xy + y^2 \geq 3xy$, so each term is at most $1/3$. This gives the bound. Equality occurs at $x = y = z$, which corresponds to $a = b = c = 1$.

This example shows a common Putnam trick: a product condition becomes a cyclic ratio condition. The substitution $a = x/y$, $b = y/z$, $c = z/x$ turns multiplicative symmetry into cyclic additive structure.

When an inequality is stuck

Try the equality case first. Then ask whether the condition is fixed sum, fixed product, or fixed norm. Fixed sum suggests convexity, Jensen, smoothing, and Cauchy. Fixed product suggests logarithms, AM-GM, or substitutions of the form $a = x/y$. Fixed norm suggests Cauchy, dot products, or geometry.

Definition 5.2 Boundary principle

For many symmetric inequalities, if equality is not at the fully equal point, the extremum occurs when variables collide or one variable reaches the boundary. This is not a theorem by itself; it is a guide to what a rigorous proof should reveal.

Example 5.4. A boundary maximum

Let $x, y, z \geq 0$ and $x + y + z = 1$. Find the maximum of $xy + yz + zx - xyz$.

Solution. Fix z and put $s = x + y = 1 - z$ and $p = xy$. Then

$$xy + yz + zx - xyz = p + zs - pz = z(1 - z) + (1 - z)p.$$

Because $0 \leq z \leq 1$, the coefficient $1 - z$ of p is nonnegative. For fixed s , the inequality $(x - y)^2 \geq 0$ gives

$$p = xy \leq \frac{(x + y)^2}{4} = \frac{(1 - z)^2}{4},$$

with equality exactly when $x = y = (1 - z)/2$. Hence the original expression is at most

$$F(z) = (1 - z) \left(z + \frac{(1 - z)^2}{4} \right) = \frac{(1 - z)(1 + z)^2}{4}.$$

On $0 \leq z \leq 1$,

$$F'(z) = \frac{(1 + z)(1 - 3z)}{4}.$$

Thus F increases up to $z = 1/3$ and decreases afterward. Its maximum is

$$F(1/3) = \frac{8}{27}.$$

Equality in both steps requires $z = 1/3$ and $x = y = (1 - z)/2 = 1/3$. Therefore the maximum is $8/27$, attained only at $x = y = z = 1/3$. The proof makes the smoothing step rigorous by fixing one variable and optimizing the other two under their fixed sum.

Example 5.5. A smoothing proof with a fixed product

Let $x, y > 0$ with fixed product $xy = p$. Show that $x + y$ is minimized when $x = y = \sqrt{p}$.

Solution. Write $x = t\sqrt{p}$ and $y = \sqrt{p}/t$. Then

$$x + y = \sqrt{p} \left(t + \frac{1}{t} \right) \geq 2\sqrt{p},$$

with equality when $t = 1$. This is AM-GM, but the substitution explains the geometry: moving away from equality in either multiplicative direction raises the sum.

The uvw mindset

For a symmetric inequality in three variables, the elementary symmetric sums

$$u = x + y + z, \quad v = xy + yz + zx, \quad w = xyz$$

are natural coordinates. If u and v are fixed, many expressions become linear in w , so the extreme cases occur when two variables are equal or at the boundary. Full uvw theory is not needed for most Putnam problems; the useful habit is to reduce symmetric expressions to symmetric sums.

Example 5.6. Two variables equal as a test

To test a symmetric inequality in x, y, z , set $y = x$ and reduce to a two-variable problem. This does not prove the result, but it often reveals the correct factor. For instance, the inequality

$$x^2 + y^2 + z^2 \geq xy + yz + zx$$

becomes $2x^2 + z^2 \geq x^2 + 2xz$, or $(x - z)^2 \geq 0$, when $x = y$. The full factorization is the sum of the three corresponding squared differences.

Problem patterns for this section

Advanced inequality problems are often optimization problems in disguise. Identify the domain, prove an extremum exists, locate possible interior equality cases, and compare boundary cases. If the variables are symmetric, test the case where two variables are equal. If a condition is multiplicative, take logarithms or use substitutions that turn products into cycles. Keep equality cases visible throughout the proof.

Proof-writing details for this section

For optimization-style inequalities, the proof must account for existence, interior behavior, and boundary behavior. A smoothing argument should say which constraint is preserved and why repeated smoothing is allowed. A tangent-line proof should say whether the function is convex or concave on the full interval being used, not merely at the equality point.

Exercises for Section 5

- (1) Prove that $\ln x \geq 1 - 1/x$ for all $x > 0$.
- (2) Let $x, y, z \geq 0$ with $x + y + z = 1$. Find the maximum of $xy + yz + zx$.
- (3) Let $a, b, c > 0$ with $abc = 1$. Prove the identity

$$a + b + c - ab - bc - ca = (a - 1)(b - 1)(c - 1),$$

and deduce that $a + b + c \geq ab + bc + ca$ holds precisely when $a = b = c = 1$ or exactly one of a, b, c is greater than 1.

- (4) Prove the exact refinement

$$x + \frac{1}{x} - 2 = \frac{(x-1)^2}{x} \quad (x > 0).$$

Deduce both $x + 1/x \geq 2$ and the local estimate $x + 1/x \geq 2 + \frac{1}{2}(x-1)^2$ for $1/2 \leq x \leq 2$.

- (5) Let $x, y, z > 0$ with $xyz = 1$. Prove that

$$\frac{x^2}{x^2 + x + 1} + \frac{y^2}{y^2 + y + 1} + \frac{z^2}{z^2 + z + 1} \geq 1.$$

- (6) Let $a_1, \dots, a_n > 0$ with fixed product. Determine when $\sum a_i$ is minimized, and prove it without naming AM-GM directly.

- (7) (**Putnam 2006 B5**) For continuous $f : [0, 1] \rightarrow \mathbb{R}$, maximize $\int_0^1 x^2 f(x) dx - \int_0^1 x(f(x))^2 dx$.

§6 Divisibility, Congruences, and Modular Arithmetic

Number theory problems become manageable when the integers are viewed through several small windows. A congruence modulo m discards most information, but it preserves exactly the information needed for divisibility by m . The art is choosing the modulus. Try small primes, powers of primes, and moduli suggested by exponents. If a problem involves squares, try modulo 4, 8, and odd primes. If it involves cubes, try modulo 7, 9, or primes p with $p - 1$ sharing factors with 3.

Definition 6.1 Congruence

For integers a, b and positive integer m , we write $a \equiv b \pmod{m}$ if m divides $a - b$.

Congruence is an equivalence relation compatible with addition and multiplication. This means that once $a \equiv b \pmod{m}$ and $c \equiv d \pmod{m}$, we may add or multiply to get $a + c \equiv b + d$ and $ac \equiv bd$ modulo m .

Theorem 6.1 Chinese remainder theorem

If m, n are relatively prime, then for any integers a, b there is an integer x satisfying

$$x \equiv a \pmod{m}, \quad x \equiv b \pmod{n}.$$

Moreover x is unique modulo mn .

Proof. Since $\gcd(m, n) = 1$, there exist integers r, s with $rm + sn = 1$. Then $x = bsn + arm$ satisfies $x \equiv a \pmod{m}$ and $x \equiv b \pmod{n}$. If two solutions exist, their difference is divisible by both m and n , hence by mn .

Example 6.1. A square obstruction

Prove that there are no integers x, y such that $x^2 + y^2 \equiv 3 \pmod{4}$.

Solution. Every integer square is congruent to 0 or 1 modulo 4. Therefore a sum of two squares is congruent to 0, 1, or 2 modulo 4, never 3. This small observation eliminates many Diophantine equations.

Definition 6.2 Order modulo m

If $\gcd(a, m) = 1$, the order of a modulo m is the least positive integer k such that $a^k \equiv 1 \pmod{m}$.

Orders turn exponential congruences into divisibility statements about exponents. For example, if $a^r \equiv a^s \pmod{m}$ and a is invertible modulo m , then $a^{r-s} \equiv 1 \pmod{m}$, so the order of a divides $r - s$.

Example 6.2. A divisibility pattern

Find all positive integers n such that 7 divides $2^n - 1$.

Solution. Compute powers of 2 modulo 7: 2, 4, 1, 2, 4, 1, ... The order of 2 modulo 7 is 3, so $2^n \equiv 1 \pmod{7}$ exactly when 3 divides n . Thus the positive solutions are $n = 3k$.

Lemma 6.1 Lifting the exponent, simple version

If p is an odd prime, $p \mid a - b$, and $p \nmid ab$, then for positive integer n ,

$$v_p(a^n - b^n) = v_p(a - b) + v_p(n).$$

Proof. Write $n = p^t m$ with $p \nmid m$. We use two elementary claims.

First, if $p \mid X - Y$, $p \nmid XY$, and $p \nmid m$, then

$$X^m - Y^m = (X - Y)(X^{m-1} + X^{m-2}Y + \cdots + Y^{m-1}).$$

Modulo p , the second factor is congruent to mY^{m-1} , which is nonzero. Therefore

$$(1) \quad v_p(X^m - Y^m) = v_p(X - Y).$$

Second, under the same conditions on X, Y ,

$$(2) \quad v_p(X^p - Y^p) = v_p(X - Y) + 1.$$

To prove this, put $d = X - Y$ and expand:

$$X^p - Y^p = (Y + d)^p - Y^p = pY^{p-1}d + \sum_{j=2}^{p-1} \binom{p}{j} Y^{p-j} d^j + d^p.$$

If $s = v_p(d) \geq 1$, the first term has valuation $s + 1$. For $2 \leq j \leq p - 1$, the prime p divides $\binom{p}{j}$, so the corresponding term has valuation at least $1 + js \geq s + 2$. The last term has valuation $ps \geq s + 2$ because p is odd. Thus the first term has strictly smaller p -adic valuation than every other term, proving (2).

Apply (1) with $X = a$ and $Y = b$ to get

$$v_p(a^m - b^m) = v_p(a - b).$$

Then apply (2) successively to the pairs

$$(a^m, b^m), (a^{mp}, b^{mp}), \dots, (a^{mp^{t-1}}, b^{mp^{t-1}}).$$

Each application raises the valuation by 1. Consequently

$$v_p(a^n - b^n) = v_p(a^m - b^m) + t = v_p(a - b) + v_p(n),$$

as claimed.

This lemma is useful but should not be overused. On the Putnam, a direct factorization and a small valuation argument are often more transparent.

Example 6.3. Choosing a modulus from the expression

Find all integer solutions to $x^2 \equiv -1 \pmod{3}$.

Solution. Squares modulo 3 are 0 and 1. The residue -1 is congruent to 2 modulo 3, so there are no solutions. The modulus was not arbitrary: the expression asks whether -1 can be a square, and modulo 3 gives the smallest obstruction.

Example 6.4. A CRT construction

Find an integer that is congruent to 1 modulo 3, congruent to 2 modulo 5, and congruent to 3 modulo 7.

Solution. Start with $x = 1 + 3a$. The condition modulo 5 gives $1 + 3a \equiv 2 \pmod{5}$, so $3a \equiv 1 \pmod{5}$ and $a \equiv 2 \pmod{5}$. Write $a = 2 + 5b$, so $x = 7 + 15b$. The condition modulo 7 gives $7 + 15b \equiv 3 \pmod{7}$, so $b \equiv 3 \pmod{7}$. Taking $b = 3$ gives $x = 52$. The solution is unique modulo 105.

A modulus selection table

Squares	Try 4, 8, and odd primes.
Cubes	Try 7 and 9, and primes where $p - 1$ interacts with 3.
Parity	Try 2 first, then 4 if signs or squares appear.
Exponents	Compute orders modulo primes or prime powers.
Factorials	Use valuations and count multiples of $p, p^2, p^3, \dots$

Example 6.5. Valuation of a factorial

Find the exponent of 5 in $100!$.

Solution. Count factors of 5. There are $\lfloor 100/5 \rfloor = 20$ multiples of 5, each contributing at least one factor. There are $\lfloor 100/25 \rfloor = 4$ multiples of 25, each contributing one additional factor. There are no multiples of 125. Hence $v_5(100!) = 24$.

Lemma 6.2 Legendre's formula

For a prime p ,

$$v_p(n!) = \left\lfloor \frac{n}{p} \right\rfloor + \left\lfloor \frac{n}{p^2} \right\rfloor + \left\lfloor \frac{n}{p^3} \right\rfloor + \cdots$$

Proof. Each multiple of p contributes one factor of p , each multiple of p^2 contributes one more, and so on. Summing these contributions counts the total number of prime factors p in the product $1 \cdot 2 \cdots n$.

Problem patterns for this section

Congruence problems reward quick experimentation. Compute residues of the main expression modulo a small modulus. If no obstruction appears, do not assume the approach failed; change the modulus. For powers, compute the order. For factorials and divisibility by high powers, use valuations. For simultaneous local conditions, solve one congruence at a time and assemble with the Chinese remainder theorem.

Proof-writing details for this section

Congruence proofs are strongest when the modulus is motivated. Once the modulus is chosen, show that every operation respects congruence. If a divisibility conclusion uses a prime power, write the valuation inequality explicitly. This is especially important when a binomial coefficient or gcd contains several overlapping prime factors.

Exercises for Section 6

- (1) Find all residues of squares modulo 8.
- (2) Prove that $x^2 + y^2 = z^2$ has no solution with x, y both odd.
- (3) Determine all n for which $5 \mid 3^n - 1$.
- (4) Prove that if $a \equiv b \pmod{m}$, then $a^k \equiv b^k \pmod{m}$ for every positive integer k .
- (5) Show that the system $x \equiv 2 \pmod{5}$, $x \equiv 3 \pmod{7}$ has a unique solution modulo 35, and find it.
- (6) Let p be an odd prime. Prove that if $p \mid a - 1$, then $p \mid 1 + a + \cdots + a^{p-1}$, and determine when p^2 divides this sum.
- (7) (**Putnam 2000 B2**) Prove that $\frac{\gcd(m, n)}{n} \binom{n}{m}$ is an integer for all integers $n \geq m \geq 1$.

§7 Diophantine Equations and Integer Methods

A Diophantine equation is an equation with integer constraints. The integer condition is not a small detail. It gives you congruences, divisibility, order, size, parity, and descent. The most useful first questions are: what are the parity restrictions, what can be bounded, and can a solution be replaced by a smaller one?

First attacks on Diophantine equations

Try parity and modular arithmetic before heavy theory. Factor if possible. If variables are positive, impose an order such as $x \leq y \leq z$. Look for a way to reduce a solution to a smaller solution. When a quadratic in one variable appears, use its second root.

Example 7.1. Pythagorean triples

Find all primitive positive integer solutions of $x^2 + y^2 = z^2$, where $\gcd(x, y, z) = 1$.

Solution. In a primitive triple, exactly one of x, y is even. Assume y is even. Then

$$y^2 = z^2 - x^2 = (z - x)(z + x).$$

The numbers $z - x$ and $z + x$ are positive even integers. Write $z + x = 2u$ and $z - x = 2v$, so $x = u - v$, $z = u + v$, and $y^2 = 4uv$. The primitivity condition implies $\gcd(u, v) = 1$. Since $uv = (y/2)^2$ and u, v are coprime, each is a square. Thus $u = m^2$, $v = n^2$, with $m > n$, giving

$$x = m^2 - n^2, \quad y = 2mn, \quad z = m^2 + n^2.$$

The conditions $\gcd(m, n) = 1$ and opposite parity are exactly what make the triple primitive.

This example shows how factorization plus coprimality converts a quadratic equation into a parametrization.

Example 7.2. Infinite descent

Prove that the equation $x^2 + y^2 = 3z^2$ has no positive integer solutions.

Solution. Suppose a solution exists with z minimal. Reducing modulo 3, we note that the only square residues are 0 and 1. Since $x^2 + y^2 \equiv 0 \pmod{3}$, we must have $x \equiv y \equiv 0 \pmod{3}$. Write $x = 3x_1$ and $y = 3y_1$. Then $9x_1^2 + 9y_1^2 = 3z^2$, so $z^2 = 3(x_1^2 + y_1^2)$, hence $3 \mid z$. Write $z = 3z_1$. Dividing by 9 gives $x_1^2 + y_1^2 = 3z_1^2$, a smaller positive solution, contradicting the minimality of z .

Definition 7.1 Vieta jumping

Vieta jumping is a descent method for equations that are quadratic in one variable. If one integer root is known, the other root is given by Vieta's formulas. One then proves the other root is a smaller nonnegative integer, contradicting minimality unless the solution has special form.

Example 7.3. A Vieta-style descent

Find all positive integer solutions of

$$x^2 + y^2 = 3xy.$$

Solution. Assume $x \leq y$. Treat the equation as a quadratic in y : $y^2 - 3xy + x^2 = 0$. If y is an integer root, the other root is $y' = 3x - y$, and $yy' = x^2$. Since $y \geq x$, we have $y' = 3x - y \leq 2x$. Also $y' = x^2/y \leq x$. If $y > x$,

then $0 < y' < x$, giving a smaller positive solution (y', x) . Infinite descent is impossible, so $x = y$. Substitution gives $2x^2 = 3x^2$, impossible for $x > 0$. Hence there are no positive integer solutions.

The argument is not about memorizing Vieta jumping. It is about seeing a quadratic structure and using the hidden second root as a descent mechanism.

Example 7.4. Bounding before solving

Find all positive integer solutions of $xy = x + y$.

Solution. Rearrange to

$$xy - x - y = 0, \quad (x - 1)(y - 1) = 1.$$

Since x, y are positive integers, both factors are nonnegative integers. Their product is 1, so $x - 1 = y - 1 = 1$, giving $x = y = 2$. The factorization is a disguised bounding argument: once the product is 1, both factors are forced.

Example 7.5. A modular obstruction with infinite families

Prove that $x^2 + y^2 = 4z + 3$ has no integer solutions.

Solution. Reduce modulo 4. A square is congruent to 0 or 1 modulo 4, so $x^2 + y^2$ is congruent to 0, 1, or 2 modulo 4. The right side is congruent to 3 modulo 4. This is impossible.

Descent template

A descent proof has four parts: assume a solution with a minimal positive measure, use the equation to build another valid solution, prove the new measure is smaller, and conclude contradiction. The measure may be a variable, a sum of variables, an absolute value, or a denominator.

Example 7.6. Descent on denominators

Prove there is no positive rational number r satisfying $r^2 = 2$.

Solution. Suppose such an $r = a/b$ exists with a, b positive integers and b minimal. Then $a^2 = 2b^2$, so a is even. Write $a = 2c$. Then $b^2 = 2c^2$, so b is even. Thus $a/b = c/(b/2)$ is the same rational number with smaller denominator, contradiction. This is the same proof of irrationality as before, now framed as a Diophantine descent.

Example 7.7. The second root check

Consider the equation $x^2 - 5xy + y^2 = 0$ in positive integers. Treating it as a quadratic in x , the product of the roots is y^2 . If one root is x , the other is $x' = 5y - x$. Any descent attempt must prove x' is positive and smaller than x . This check is the heart of Vieta jumping. Without it, the second root is only algebra, not a proof.

Problem patterns for this section

For Diophantine equations, the first goal is not to solve but to reduce. Order the variables, factor the equation, and test small moduli. If a variable appears quadratically, inspect the second root. If a solution can be scaled down, descent may prove impossibility. If the variables are positive, inequalities can bound one variable in terms of another and make a finite search rigorous.

Proof-writing details for this section

For Diophantine equations, a complete solution has a necessary part and a sufficient part. Modular restrictions, size bounds, and descent arguments usually give necessity. The final line must still verify every remaining candidate. A common mistake is to find a plausible solution and forget to exclude all other residue classes or boundary cases.

Exercises for Section 7

- (1) Find all integer solutions to $x^2 - y^2 = 15$.
- (2) Prove that $x^2 + y^2 = 7$ has no integer solutions.
- (3) Find all positive integer solutions of $1/x + 1/y = 1/6$.
- (4) Prove that if $x^2 = 2y^2$ in integers, then $x = y = 0$.
- (5) Find all positive integer solutions of $x^2 - y^2 = 2z^2$ with x, y consecutive.
- (6) Let x, y be positive integers such that $xy + 1$ divides $x^2 + y^2$, and write

$$k = \frac{x^2 + y^2}{xy + 1}.$$

Prove, by a Vieta descent on the larger of x and y , that k is a perfect square.

- (7) **(Putnam 2001 A5)** Find the unique positive integers a, n such that $a^{n+1} - (a+1)^n = 2001$.

§8 Arithmetic Functions and Prime Factorization Techniques

Prime factorization is the coordinate system of the positive integers. Questions about divisibility become questions about exponents of primes. Arithmetic functions such as $\varphi(n)$, $d(n)$, and $\sigma(n)$ are useful because they translate multiplicative structure into formulas.

Definition 8.1 Valuation

For a prime p and a nonzero integer n , $v_p(n)$ is the exponent of p in the prime factorization of n .

The valuation turns multiplication into addition: $v_p(ab) = v_p(a) + v_p(b)$. It turns divisibility into comparison: $a \mid b$ if and only if $v_p(a) \leq v_p(b)$ for every prime p .

Example 8.1. Counting divisors

If $n = p_1^{a_1} \cdots p_k^{a_k}$, prove that the number of positive divisors of n is

$$d(n) = (a_1 + 1) \cdots (a_k + 1).$$

Solution. A positive divisor has the form $p_1^{b_1} \cdots p_k^{b_k}$ with $0 \leq b_i \leq a_i$. For each i , there are $a_i + 1$ choices for b_i , and the choices are independent. Multiplying the numbers of choices gives the formula.

Theorem 8.1 Euler's theorem

If $\gcd(a, n) = 1$, then

$$a^{\varphi(n)} \equiv 1 \pmod{n}.$$

Proof. Let $r_1, \dots, r_{\varphi(n)}$ be the reduced residue classes modulo n . Multiplication by a permutes these classes because a is invertible modulo n . Therefore

$$(ar_1)(ar_2) \cdots (ar_{\varphi(n)}) \equiv r_1 r_2 \cdots r_{\varphi(n)} \pmod{n}.$$

The product $r_1 \cdots r_{\varphi(n)}$ is invertible modulo n , so it may be canceled, leaving $a^{\varphi(n)} \equiv 1 \pmod{n}$.

Example 8.2. A quick exponent reduction

Find the remainder when 7^{100} is divided by 10.

Solution. Since $\varphi(10) = 4$ and $\gcd(7, 10) = 1$, Euler's theorem gives $7^4 \equiv 1 \pmod{10}$. Since 100 is a multiple of 4, we get $7^{100} \equiv 1 \pmod{10}$.

Definition 8.2 Multiplicative function

A function $f : \mathbb{N} \rightarrow \mathbb{C}$ is multiplicative if $f(mn) = f(m)f(n)$ whenever $\gcd(m, n) = 1$.

Many divisor sums are multiplicative because prime-power choices are independent. This is the same product principle from counting, now applied to factorization.

Lemma 8.1 Divisor sum formula

If $n = p_1^{a_1} \cdots p_k^{a_k}$, then

$$\sum_{d \mid n} d = \prod_{i=1}^k \frac{p_i^{a_i+1} - 1}{p_i - 1}.$$

Proof. Each divisor is obtained by choosing one term from each list $1, p_i, p_i^2, \dots, p_i^{a_i}$. Expanding the product of these finite geometric sums selects each divisor exactly once.

Example 8.3. Euler phi from prime powers

Compute $\varphi(72)$.

Solution. Since $72 = 2^3 \cdot 3^2$, use multiplicativity and the prime-power formula:

$$\varphi(72) = \varphi(2^3)\varphi(3^2) = (8 - 4)(9 - 3) = 4 \cdot 6 = 24.$$

The formula counts integers not divisible by the relevant prime in each prime-power modulus.

Lemma 8.2 Multiplicativity of Euler phi

If $\gcd(m, n) = 1$, then $\varphi(mn) = \varphi(m)\varphi(n)$.

Proof. By the Chinese remainder theorem, each residue modulo mn corresponds uniquely to a pair of residues modulo m and modulo n . A residue is relatively prime to mn exactly when its two components are relatively prime to m and to n . Hence the number of reduced residues modulo mn is the product of the two counts.

Example 8.4. A divisor product identity

If n has $d(n)$ positive divisors, prove that the product of all positive divisors of n is $n^{d(n)/2}$.

Solution. Pair each divisor d with n/d . Each pair has product n . If $d \neq n/d$, this is a two-element pair. If $d = n/d$, then n is a square and the middle divisor contributes $\sqrt{n}$, which is exactly what the formula gives because $d(n)$ is then odd. More uniformly, multiplying all equations $d \cdot (n/d) = n$ over all divisors counts each divisor twice, so the square of the desired product is $n^{d(n)}$.

Definition 8.3 Mobius inversion, light version

The Mobius function $\mu(n)$ is 0 if n is divisible by a square prime factor, and otherwise $(-1)^k$ if n is a product of k distinct primes. It satisfies

$$\sum_{d|n} \mu(d) = \begin{cases} 1, & n = 1, \\ 0, & n > 1. \end{cases}$$

Mobius inversion is not needed often, but the identity above is a useful inclusion-exclusion over prime divisors. It says that alternating over squarefree divisors filters out everything except 1.

Example 8.5. Counting coprime integers by inclusion-exclusion

The number of integers $1 \leq k \leq n$ relatively prime to n can be counted by excluding multiples of the prime divisors of n . If n has distinct prime divisors $p_1, \dots, p_r$, then

$$\varphi(n) = n \prod_{i=1}^r \left(1 - \frac{1}{p_i}\right).$$

This formula follows from inclusion-exclusion and is often faster than listing residues.

Problem patterns for this section

Arithmetic-function problems become easier when prime powers are handled first. Prove the formula for p^a , then use multiplicativity if coprime factors separate. When a divisor sum appears, think of choosing exponents independently. When a quotient such as $\sigma(n)/n$ appears, rewrite it as a product over primes. When a square condition appears, translate it into parity of valuations.

Proof-writing details for this section

Arithmetic-function proofs usually become clean after prime factorization. When proving an identity involving divisors, lcm, gcd, or factorials, compare the exponent of an arbitrary prime p on both sides. This converts a global multiplicative statement into a counting statement about how many times powers of p occur.

Exercises for Section 8

- (1) Compute $d(360)$ and $\sum_{d|360} d$.
- (2) Prove that $d(n)$ is odd if and only if n is a perfect square.
- (3) Prove that $\varphi(p^a) = p^a - p^{a-1}$ for prime p .
- (4) Show that if m, n are coprime, then $d(mn) = d(m)d(n)$.
- (5) Find all n such that $\varphi(n) = n/2$.
- (6) Prove that if $\sum_{d|n} d = 2n$, then n cannot be prime. Verify directly that 6, 28, and 496 satisfy the equation.
- (7) **(Putnam 2003 B3)** Prove that $n! = \prod_{i=1}^n \text{lcm}\{1, 2, \dots, \lfloor n/i \rfloor\}$ for every positive integer n .

§9 Counting Methods and Double Counting

Counting is the art of assigning names to the same collection in two different ways. A good double count does not feel like a trick. It identifies a set of pairs, triples, paths, subsets, or incidences that naturally has two descriptions. The Putnam often hides a counting problem inside algebraic notation or a probability question.

Build the counted set

When you see a sum, ask what each term counts. When you see a product, ask what independent choices it represents. When you see an identity, try to find one set counted by both sides. When you see an average, count incidences and divide.

Example 9.1. Handshakes as double counting

In a finite graph with vertex set V and edge set E , prove that

$$\sum_{v \in V} \deg(v) = 2|E|.$$

Solution. Count ordered incidences (v, e) where vertex v lies on edge e . For each vertex v , there are $\deg(v)$ incident edges, giving the left side. For each edge, there are exactly two endpoints, giving the right side. The two counts refer to the same incidence set.

Example 9.2. A binomial sum with a marker

Prove

$$\sum_{k=0}^n k \binom{n}{k} = n2^{n-1}.$$

Solution. Count pairs (S, s) where $S \subseteq \{1, \dots, n\}$ and $s \in S$. If $|S| = k$, there are $k \binom{n}{k}$ such pairs, giving the left side. Alternatively, choose the marked element s in n ways and choose freely whether each of the remaining $n - 1$ elements is in S , giving $n2^{n-1}$.

Theorem 9.1 Inclusion-exclusion

For finite sets $A_1, \dots, A_n$,

$$\left| \bigcup_{i=1}^n A_i \right| = \sum_i |A_i| - \sum_{i < j} |A_i \cap A_j| + \sum_{i < j < k} |A_i \cap A_j \cap A_k| - \dots$$

Proof. Fix an element contained in exactly r of the sets. Its total contribution to the right side is

$$\binom{r}{1} - \binom{r}{2} + \binom{r}{3} - \dots + (-1)^{r+1} \binom{r}{r} = 1,$$

by the binomial theorem applied to $(1 - 1)^r$. Elements contained in no set contribute zero. Thus every element of the union is counted exactly once.

Example 9.3. Derangements

Let D_n be the number of permutations of $\{1, \dots, n\}$ with no fixed point. Show that

$$D_n = n! \sum_{k=0}^n \frac{(-1)^k}{k!}.$$

Solution. Let A_i be the set of permutations fixing i . We want to count permutations outside $A_1 \cup \cdots \cup A_n$. If a specified set of k positions is fixed, the remaining $n - k$ elements may be permuted arbitrarily, giving $(n - k)!$ permutations. Inclusion-exclusion gives

$$D_n = \sum_{k=0}^n (-1)^k \binom{n}{k} (n - k)! = n! \sum_{k=0}^n \frac{(-1)^k}{k!}.$$

Example 9.4. Weighted double counting

Let $x_1, \dots, x_n$ be real numbers. Prove that

$$\sum_{1 \leq i < j \leq n} (x_i + x_j) = (n - 1) \sum_{i=1}^n x_i.$$

Solution. Each x_i appears in exactly $n - 1$ pairs (i, j) with $i < j$ or $j < i$. Therefore its total contribution to the left side is $(n - 1)x_i$. Summing over i gives the identity. This is double counting with weights rather than objects.

Example 9.5. Counting lattice paths

How many shortest paths go from $(0, 0)$ to (a, b) using only steps right and up?

Solution. Any shortest path has exactly a right steps and b up steps, in some order. Thus it is determined by the positions of the a right steps among the $a + b$ total steps. The number is

$$\binom{a + b}{a} = \binom{a + b}{b}.$$

This model is a source of many binomial identities because a path can be decomposed by its first visit to a line, a point, or a boundary.

Example 9.6. Vandermonde's identity

Prove that

$$\sum_{k=0}^r \binom{m}{k} \binom{n}{r - k} = \binom{m + n}{r}.$$

Solution. Choose r people from a group split into two rooms of sizes m and n . If k chosen people come from the first room, then $r - k$ come from the second. Summing over k gives the left side, while choosing directly gives the right side.

Signs in inclusion-exclusion

The sign pattern is not a formula to memorize. An object lying in exactly r forbidden sets contributes

$$\binom{r}{1} - \binom{r}{2} + \binom{r}{3} - \cdots + (-1)^{r+1} \binom{r}{r} = 1$$

when counting the union. For the complement, the contribution becomes

$$1 - \binom{r}{1} + \binom{r}{2} - \cdots + (-1)^r \binom{r}{r},$$

which is 0 unless $r = 0$. This is the whole mechanism.

Example 9.7. A recurrence from counting the last step

Let a_n be the number of binary strings of length n with no two consecutive 1s. Then $a_n = a_{n-1} + a_{n-2}$. A valid string either ends in 0, preceded by any valid string of length $n - 1$, or ends in 10, preceded by any valid string of length $n - 2$. This last-step decomposition is the combinatorial version of induction.

Problem patterns for this section

Counting problems become clear when the counted objects are named. If a proof starts with a sum, define the set whose elements contribute to that sum. If it starts with an identity, describe the same set in two ways. If it involves avoiding restrictions, use inclusion-exclusion or count the complement. If it involves paths or strings, classify by the first step, last step, or first time a boundary is reached.

Proof-writing details for this section

Counting proofs should make the sample space and the counted objects explicit. If a sum is evaluated by double counting, describe the set counted by both sides. If expectation is used, define the indicator variables and justify their probabilities. Linearity of expectation does not require independence, and saying so often shortens the proof.

Exercises for Section 9

- (1) Count the number of subsets of $\{1, \dots, n\}$ in two ways to prove $\sum_k \binom{n}{k} = 2^n$.
- (2) Prove $\sum_k \binom{n}{k}^2 = \binom{2n}{n}$ by counting subsets.
- (3) How many positive integers at most 1000 are divisible by 2, 3, or 5?
- (4) Count the number of onto functions from an n -element set to a 3-element set.
- (5) Let $A_1, \dots, A_m$ be subsets of an n -element set, each of size r . If every element lies in exactly s of the subsets, prove that $mr = ns$.
- (6) Let S be a set of n points in the plane with no three collinear. Count ordered triples (P, ℓ, Q) such that $P, Q \in S$, the line ℓ is determined by two points of S , $P \in \ell$, and $Q \notin \ell$.
- (7) (**Putnam 2006 A4**) Let $n \geq 2$. For a uniformly chosen permutation of $\{1, \dots, n\}$, determine the average number of local maxima, where endpoints are treated as local maxima when they exceed their only neighbor.

§10 Pigeonhole Principle, Extremal Principles, and Probabilistic Thinking

The pigeonhole principle is the beginning of existence theory. It says that if too many objects are placed into too few boxes, structure appears. Extremal principles say that a finite nonempty collection has a largest or smallest object, and that this object is constrained by its extremality. Probabilistic thinking says that if the average value is large, some example has value at least that average.

Theorem 10.1 Pigeonhole principle

If N objects are placed into k boxes, then some box contains at least $\lceil N/k \rceil$ objects.

Proof. If every box contained at most $\lceil N/k \rceil - 1$ objects, the total number of objects would be at most $k(\lceil N/k \rceil - 1) < N$, a contradiction.

Example 10.1. Subset sums modulo n

Let $a_1, \dots, a_n$ be integers. Prove that some nonempty consecutive block has sum divisible by n .

Solution. Consider partial sums $s_k = a_1 + \dots + a_k$ for $1 \leq k \leq n$. If some $s_k \equiv 0 \pmod{n}$, then $a_1 + \dots + a_k$ is the desired block. Otherwise the n sums occupy only the $n - 1$ nonzero residue classes modulo n , so two of them are congruent. If $s_i \equiv s_j$ with $i < j$, then $a_{i+1} + \dots + a_j = s_j - s_i$ is divisible by n .

Definition 10.1 Extremal argument

An extremal argument selects an object with maximal or minimal value of a carefully chosen quantity, then proves that this object cannot be improved, moved, deleted, or replaced unless the desired conclusion holds.

Example 10.2. A closest pair argument

A finite set of points in the plane has the property that the midpoint of any two points in the set also belongs to the set. Prove that the set has at most one point.

Solution. Suppose there are at least two points. Among all pairs of distinct points in the finite set, choose A, B with minimal positive distance. By hypothesis, the midpoint M of AB also lies in the set. Since $A \neq B$, the point M is distinct from both A and B , and $AM = BM = AB/2$. This is a smaller positive distance between two points of the set, contradicting the choice of A, B . Hence the set has at most one point.

This example reveals the purpose of an extremal choice. The chosen pair is not special because we can compute it. It is special because any new object forced by the hypothesis cannot be smaller than it.

Theorem 10.2 First moment principle

If a finite collection of real numbers has average A , then at least one number is at least A and at least one number is at most A .

Proof. If every number were less than A , the average would be less than A . The other direction is identical.

Example 10.3. A probabilistic construction

Show that every graph with m edges has a bipartition of its vertices that cuts at least $m/2$ edges.

Solution. Put each vertex independently into one of two parts with probability $1/2$. For any fixed edge, the probability that its endpoints land in different parts is $1/2$. Therefore the expected number of cut edges is $m/2$. Since the average over all random partitions is $m/2$, some partition cuts at least $m/2$ edges.

Theorem 10.3 Erdos-Szekeres monotone subsequence lemma

Among more than rs distinct real numbers, there is an increasing subsequence of length $r + 1$ or a decreasing subsequence of length $s + 1$.

Proof. To each term assign the pair (a, b) , where a is the length of the longest increasing subsequence starting at that term and b is the length of the longest decreasing subsequence starting at that term. If two terms have the same pair and the earlier term is smaller, then its increasing length would be larger than the later term's increasing length. If the earlier term is larger, its decreasing length would be larger. Thus all pairs are distinct. If all increasing lengths are at most r and all decreasing lengths are at most s , there are at most rs possible pairs, contradiction.

Example 10.4. A direct application

Prove that any sequence of 10 distinct real numbers contains an increasing subsequence of length 4 or a decreasing subsequence of length 4.

Solution. Use the lemma with $r = s = 3$. Since $10 > 9$, the sequence contains an increasing subsequence of length 4 or a decreasing subsequence of length 4.

Example 10.5. Averaging over choices

A set of n real numbers has average A . Prove that there is a subset of size k whose average is at least A .

Solution. Choose a k -element subset uniformly at random. Each original number appears in the random subset with probability k/n . Therefore the expected sum of the chosen subset is $(k/n)(x_1 + \cdots + x_n) = kA$, so the expected average is A . At least one subset has average at least the expected average.

Probabilistic method without probability language

The same proof can be written as double counting. Sum the averages of all k -element subsets. Each x_i appears in exactly $\binom{n-1}{k-1}$ subsets, so the average of the subset averages is A . Probability is often just compact notation for a double count.

Example 10.6. Extremal deletion

Suppose a graph has average degree at least d . Delete vertices of degree less than $d/2$ as long as possible. If all vertices are deleted, count the number of removed edges and get a contradiction to the original average. Therefore a nonempty subgraph remains with minimum degree at least $d/2$. This method is a standard way to turn an average condition into a local minimum condition.

Problem patterns for this section

Pigeonhole arguments require boxes. If the boxes are not obvious, create a signature: residue class, parity vector, interval, degree, or pair of subsequence lengths. Extremal arguments require a measure. If the first measure does not lead to contradiction, change it. Probabilistic arguments require a random experiment

with an expected value that is easy to compute. The experiment should be simpler than the object you seek.

Proof-writing details for this section

Pigeonhole and extremal arguments require clear boxes and clear objects. The proof should name the boxes, count the objects, and explain what collision or extremal choice follows. For averaging arguments, state the average and why at least one object is at least or at most that average.

Exercises for Section 10

- (1) Among any 13 people, prove that two have birthdays in the same month.
- (2) Prove that among any $n + 1$ integers, two have difference divisible by n .
- (3) Let $a_1, \dots, a_n$ be real numbers. Prove that some a_i is at least the average.
- (4) Prove that every tournament has a vertex with outdegree at least $(n - 1)/2$.
- (5) Show that every set of $n + 1$ integers chosen from $\{1, 2, \dots, 2n\}$ contains two numbers such that one divides the other.
- (6) Let $m \geq 1$. If $m^2 + 1$ points lie in a unit square, prove that two of them are at distance at most $\sqrt{2}/m$.
- (7) (**Putnam 2002 A2**) Given any five points on a sphere, prove that some four of them lie on a closed hemisphere.

§11 Recurrences, Generating Functions, and Discrete Processes

Recurrences describe objects built from smaller objects. Generating functions turn recurrences into algebra. Discrete processes often hide a recurrence in time, a state graph, or an invariant. A Putnam recurrence problem is often not about solving a recurrence exactly. It may ask for divisibility, monotonicity, periodicity, or asymptotic size.

Definition 11.1 Ordinary generating function

For a sequence $(a_n)_{n \geq 0}$, its ordinary generating function is

$$A(x) = \sum_{n \geq 0} a_n x^n.$$

The coefficient of x^n stores the n th term. Multiplication of generating functions corresponds to convolution, which is why generating functions are natural for counting decompositions.

Example 11.1. Fibonacci recurrence

Let $F_0 = 0$, $F_1 = 1$, and $F_{n+2} = F_{n+1} + F_n$. Find the generating function $F(x) = \sum_{n \geq 0} F_n x^n$.

Solution. Multiply the recurrence by x^{n+2} and sum for $n \geq 0$:

$$\sum_{n \geq 0} F_{n+2} x^{n+2} = \sum_{n \geq 0} F_{n+1} x^{n+2} + \sum_{n \geq 0} F_n x^{n+2}.$$

The left side is $F(x) - x$, while the right side is $xF(x) + x^2F(x)$. Hence

$$F(x) - x = xF(x) + x^2F(x),$$

so

$$F(x) = \frac{x}{1 - x - x^2}.$$

Theorem 11.1 Linear recurrence principle

A sequence satisfying a linear recurrence with constant coefficients is controlled by the roots of its characteristic polynomial. If the roots are distinct, the sequence is a linear combination of powers of those roots.

Proof idea. The shift operator sends a_n to a_{n+1} . A recurrence such as $a_{n+2} = pa_{n+1} + qa_n$ says that the sequence is killed by a polynomial in the shift operator. Sequences of the form r^n are eigenvectors of this operator, and substituting r^n gives the characteristic equation.

Example 11.2. A parity pattern

Prove that the Fibonacci number F_n is even if and only if 3 divides n .

Solution. Work modulo 2. The recurrence produces

$$F_0, F_1, F_2, F_3, F_4, F_5, \dots \equiv 0, 1, 1, 0, 1, 1, \dots \pmod{2}.$$

Once the pair (F_n, F_{n+1}) repeats, the entire sequence repeats. The pair $(0, 1)$ occurs at $n = 0$ and again at $n = 3$, so the period is 3. Thus $F_n \equiv 0 \pmod{2}$ exactly when $n \equiv 0 \pmod{3}$.

Example 11.3. Generating functions for coin choices

How many ways are there to write n as a sum of 1s and 2s, where order does not matter?

Solution. The number of 1s can be any nonnegative integer, and the number of 2s can be any nonnegative integer. The generating function is

$$(1 + x + x^2 + \cdots)(1 + x^2 + x^4 + \cdots) = \frac{1}{(1-x)(1-x^2)}.$$

The coefficient of x^n counts the number of choices of $b \geq 0$ such that $2b \leq n$, so it is $\lfloor n/2 \rfloor + 1$.

Example 11.4. State vectors for a recurrence

Let $a_{n+2} = a_{n+1} + a_n$. Show that

$$\begin{pmatrix} a_{n+1} \\ a_n \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}^n \begin{pmatrix} a_1 \\ a_0 \end{pmatrix}.$$

Solution. The recurrence is equivalent to

$$\begin{pmatrix} a_{n+2} \\ a_{n+1} \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} a_{n+1} \\ a_n \end{pmatrix}.$$

Iterating this matrix identity gives the formula. Matrix form is useful when a recurrence involves several previous terms or when modular periodicity is the target.

Example 11.5. A generating function product

Find the generating function for partitions into distinct parts, where each positive integer part may be used at most once.

Solution. For each part size k , the choice is either to use it or not use it, contributing a factor $1 + x^k$. Multiplying independent choices gives

$$\prod_{k \geq 1} (1 + x^k).$$

The coefficient of x^n counts partitions of n into distinct parts. This expression is formal; for any fixed coefficient, only finitely many factors matter.

Finite state processes

If a deterministic process has finitely many states, then every trajectory is eventually periodic. If the transition rule is reversible, then every trajectory is purely periodic from the start. This principle is invaluable for recurrences modulo m , games with finite positions, and automata-like problems.

Example 11.6. A recurrence modulo m

Let $a_{n+2} = a_{n+1} + a_n$ modulo m . The state at time n is the pair (a_n, a_{n+1}) , and there are only m^2 possible states. Hence some state repeats. Because the map $(u, v) \mapsto (v, u + v)$ is reversible modulo m , with inverse $(v, w) \mapsto (w - v, v)$, the sequence of states is periodic from the beginning.

Example 11.7. Solving by characteristic roots

For $a_{n+2} = 5a_{n+1} - 6a_n$, try $a_n = r^n$. This gives $r^2 = 5r - 6$, so $r = 2$ or 3 . Therefore $a_n = A2^n + B3^n$ for constants A, B determined by a_0, a_1 . Even when exact formulas are unnecessary, the roots tell growth rate and possible divisibility behavior.

Problem patterns for this section

Recurrence problems ask for the right state. A one-term recurrence has state a_n ; a two-term recurrence has state (a_n, a_{n+1}) ; a process with memory needs all data required to determine the next move. For counting recurrences, decompose by the last step. For modular recurrences, count states to prove periodicity. For exact formulas, use generating functions or characteristic roots.

Proof-writing details for this section

Recurrence and generating-function proofs should include initial conditions. A closed form is not proved until it satisfies both the recurrence and the starting values. For generating functions, every coefficient extraction must respect the range of summation, especially when indices are shifted.

Exercises for Section 11


- (1) Solve $a_{n+1} = 2a_n$ with $a_0 = 3$.
- (2) Find the generating function for $a_n = n$.
- (3) Prove that $F_{n+1}F_{n-1} - F_n^2 = (-1)^n$.
- (4) Find the number of ways to write n as a sum of 1s, 3s, and 4s where order matters.
- (5) Let $a_{n+2} = 3a_{n+1} - 2a_n$, $a_0 = 0$, $a_1 = 1$. Find a closed form.
- (6) Let R be a finite set and let $x_{n+2} = F(x_n, x_{n+1})$ for a fixed map $F : R^2 \rightarrow R$. Prove that the state sequence (x_n, x_{n+1}) is eventually periodic. Prove that it is periodic from the initial state whenever the state map $T(u, v) = (v, F(u, v))$ is bijective.
- (7) (**Putnam 1999 A3**) If $\frac{1}{1-2x-x^2} = \sum_{n \geq 0} a_n x^n$, prove that for each $n \geq 0$ there is an integer m such that $a_n^2 + a_{n+1}^2 = a_m$.

§12 Graph Theory for the Putnam


Graphs encode relations. Vertices are objects, edges are pairwise relationships, directions encode choices, colors encode classes, and weights encode costs. Many Putnam problems that do not mention graphs become easy after drawing one. The graph method is especially strong when the problem has words like connected, pair, relation, reachable, compatible, adjacent, or process.

Definition 12.1 Graph basics

A graph $G = (V, E)$ consists of vertices V and edges E , where each edge joins two vertices. A path is a sequence of adjacent vertices. A graph is connected if every two vertices are joined by a path. A tree is a connected graph with no cycles.



A graph with a cycle



A tree

Theorem 12.1 Tree edge count

A finite connected graph with n vertices is a tree if and only if it has $n - 1$ edges and no redundant edge.

Proof. A tree can be built from one vertex by repeatedly adding one new vertex attached by one edge, so a tree with n vertices has $n - 1$ edges. Conversely, a connected graph with a cycle has an edge on the cycle whose removal keeps the graph connected. Repeating this removal until no cycles remain gives a connected acyclic graph, a tree, with $n - 1$ edges. Therefore a connected graph with $n - 1$ edges could not have had a cycle.

Example 12.1. Odd degree toll roads

A connected graph has an even number of vertices. Prove that one can choose some edges so that every vertex is incident to an odd number of chosen edges.

Solution. It is enough to solve the problem on a spanning tree, since chosen tree edges are also graph edges. Root the tree at any vertex. Work from the leaves upward. For each nonroot vertex, decide whether to choose the edge joining it to its parent so that the vertex ends with odd chosen degree after all edges to its children have already been decided. This determines every parent edge. At the end, every nonroot vertex has odd chosen degree. The root also has odd chosen degree because the sum of all chosen degrees is twice the number of chosen edges, hence even, and an even number of vertices are involved. If all nonroot vertices have odd degree and there are an odd number of nonroot vertices, the root must also have odd degree.

Theorem 12.2 Euler's formula for planar graphs

If a connected planar graph has V vertices, E edges, and F faces, then

$$V - E + F = 2.$$

Proof. If the graph is a tree, then $E = V - 1$ and there is one face, so the formula holds. If the graph has a cycle, remove an edge on a cycle. This keeps the graph connected and merges two faces into one, so both E and F decrease by 1. The quantity $V - E + F$ is unchanged. Repeating until a tree remains proves the formula.

Example 12.2. A planar bound

Prove that a simple planar graph with $V \geq 3$ has at most $3V - 6$ edges.

Solution. In a simple planar graph, every face is bounded by at least three edges, and every edge is counted twice along face boundaries. Thus $3F \leq 2E$. Euler's formula gives $F = 2 - V + E$, so $3(2 - V + E) \leq 2E$. Rearranging yields $E \leq 3V - 6$.

Theorem 12.3 Bipartite graph criterion

A graph is bipartite if and only if it has no cycle of odd length.

Proof. A bipartite graph has every cycle alternating between the two parts, so every cycle has even length. Conversely, suppose a connected graph has no odd cycle. Choose a root vertex and put each vertex into one of two classes according to whether its distance from the root is even or odd. If an edge joined two vertices in the same class, the two root paths plus that edge would contain an odd cycle after deleting any common initial part. Hence every edge goes between the two classes. Apply the argument to each connected component.

Example 12.3. Encoding compatibility

A group of people each either likes or dislikes every other person. Suppose there is no group of three people who pairwise like each other and no group of three who pairwise dislike each other. What graph statement is this?

Solution. Make a graph whose vertices are people and whose edges join pairs who like each other. The condition says the graph has no triangle and its complement has no triangle. This is the language of Ramsey theory. The famous statement $R(3, 3) = 6$ says that among six people, one of those two triangles must exist.

Lemma 12.1 Hall's marriage theorem

A bipartite graph with left vertex set L has a matching covering all vertices of L if and only if for every subset $S \subseteq L$, the neighbor set $N(S)$ has size at least $|S|$.

Hall's theorem is a powerful existence criterion. On the Putnam, it is often enough to know when it might apply: assignments, systems of distinct representatives, matching tasks to resources, and choosing distinct objects from overlapping sets.

Example 12.4. Distinct representatives

Let $A_1, \dots, A_n$ be finite sets. A system of distinct representatives is a choice of $a_i \in A_i$ with all a_i distinct. Hall's theorem says such a system exists exactly when every union of k of the sets contains at least k elements. Build a bipartite graph with set-vertices on the left and element-vertices on the right, joining A_i to the elements it contains.

Graph encoding triggers

Use graphs when a problem has pairwise relations, reachability, local degree restrictions, parity of incidences, or choices that must be distinct. Use directed graphs when choices have direction, dependency, or a next step. Use edge colors when two relations exist on the same set.

Problem patterns for this section

Graph problems often begin by deleting or contracting. Trees invite induction on leaves. Degree conditions invite the handshaking lemma or averaging. Planarity invites Euler's formula. Matching problems invite Hall's theorem. Coloring problems invite parity, odd cycles, or greedy algorithms. When a non-graph problem has pairwise relations, create the graph yourself and let degrees, paths, and components organize the information.

Proof-writing details for this section

Graph proofs should translate the original language into vertices, edges, paths, cuts, or matchings. After the translation, prove the graph statement with its hypotheses stated. If Hall's theorem is used, the proof must verify Hall's condition for an arbitrary set, not just for a representative example.

Exercises for Section 12


- (1) Prove the handshaking lemma.
- (2) Show that every finite tree with at least two vertices has at least two leaves.
- (3) Prove that every connected graph contains a spanning tree.
- (4) Use Euler's formula to prove that K_5 is not planar.
- (5) Prove that every graph with average degree greater than $2k$ contains a subgraph with minimum degree at least $k + 1$.
- (6) In a tournament, prove that there is a directed path visiting every vertex exactly once.
- (7) (**Putnam 2012 B3**) In a round-robin tournament of $2n$ teams played over $2n - 1$ days, with every team playing once each day, prove that one can choose one winning team from each day without choosing any team twice.

§13 Euclidean Geometry and Transformational Methods

Geometry on the Putnam is often solved by changing the language. Synthetic geometry uses angles, circles, similarity, and power of a point. Transformational geometry uses rotations, reflections, homotheties, and inversion. Analytic geometry uses coordinates. The best method is the one that makes the invariant visible.

Definition 13.1 Power of a point

For a circle with center O and radius r , the power of a point P is $PO^2 - r^2$. If a line through P meets the circle at A and B , then $PA \cdot PB$ equals the power of P , with directed lengths.



Theorem 13.1 Intersecting secants

If two lines through a point P meet a circle at A, B and C, D , then

$$PA \cdot PB = PC \cdot PD.$$

Proof. The triangles formed by the two secants are similar by equal angles subtending the same arcs. The proportional sides give $PA/PC = PD/PB$, which rearranges to the stated equality.

Example 13.1. A length from power of a point

A point P outside a circle has a tangent segment PT of length 6. A secant through P meets the circle at A and B , with $PA = 4$. Find PB .

Solution. The tangent form of power of a point gives $PT^2 = PA \cdot PB$. Thus $36 = 4 \cdot PB$, so $PB = 9$.

Definition 13.2 Similarity as a coordinate-free scale

Two figures are similar if one can be obtained from the other by translations, rotations, reflections, and uniform scaling. Similarity preserves angles and ratios of corresponding lengths.

Example 13.2. A rotation trick

Let ABC be an equilateral triangle and let P be a point inside it. Show that the three segments PA , PB , and PC can be arranged to form a triangle exactly when no one length exceeds the sum of the other two.

Solution. Rotate the plane by 60° about A , sending B to C . The point P moves to P' . Then $AP = AP'$ and $\angle PAP' = 60^\circ$, so $PP' = AP$. Also $P'C = PB$. Thus the broken path $P - P' - C$ has side lengths PA , PB , and PC in a triangle-like configuration. The triangle inequality condition is precisely the condition for these lengths to close. The rotation converts three distances from one moving point into sides of a triangle.

Angle chasing discipline

Do not chase every angle. Mark the target angle first, then mark only angles connected to it by cyclic quadrilaterals, parallel lines, isosceles triangles, or tangent-radius perpendicularity. A clean angle chase is a directed path, not a filled diagram.

Example 13.3. Cyclic quadrilateral test

Let A, B, C, D be points with $\angle ABC = \angle ADC$. Prove that A, B, C, D lie on a common circle, assuming the points are positioned so the angles subtend the same chord AC .

Solution. The set of points X from which segment AC is seen under a fixed angle is an arc of a circle through A and C . Since both B and D see chord AC under the same angle, they lie on the same circle through A and C . This is the converse of the inscribed angle theorem.

Example 13.4. Inversion intuition

Inversion in a circle centered at O with radius R sends a point P to the point P' on ray OP satisfying $OP \cdot OP' = R^2$. Lines and circles through O change into lines or circles in simple ways. You do not need inversion for most Putnam geometry, but the idea is useful: replace products of distances by reciprocal distances from a fixed center.

Example 13.5. Area as a hidden invariant

Let D lie on side BC of triangle ABC . Prove that

$$\frac{[ABD]}{[ACD]} = \frac{BD}{DC}.$$

Solution. The two triangles have bases BD and DC on the same line and share the same altitude from A to line BC . Since area is one half base times height, the ratio of areas is the ratio of the bases. This simple fact turns many cevian problems into algebraic ratios.

Theorem 13.2 Ceva's theorem

For points $D \in BC$, $E \in CA$, and $F \in AB$, the cevians AD , BE , and CF are concurrent if and only if

$$\frac{BD}{DC} \cdot \frac{CE}{EA} \cdot \frac{AF}{FB} = 1$$

using directed lengths.

Proof. If the cevians meet at P , compare areas of the six small triangles around P . Ratios of areas sharing an altitude give the three side ratios, and the products telescope to 1. Conversely, if the product condition holds, let AD and BE meet at P and let CP meet AB at F' . The already proved direction gives the same ratio for F' , so $F' = F$.

Problem patterns for this section

Geometry problems should be redrawn several times. Mark equal angles, equal lengths, cyclic quadrilaterals, and parallel lines. If the diagram has many ratios, use area or mass points. If it has circles and products of lengths, use power of a point. If it has a special angle such as 60° or 90° , consider a rotation or reflection. If all else fails, choose coordinates that respect the symmetry.

Proof-writing details for this section

Geometry proofs become readable when the construction is fixed before the computation begins. Name the points, state which triangles are similar or cyclic, and identify which lengths are equal. A diagram helps intuition, but the proof must still say which theorem justifies each relation.

Exercises for Section 13

- (1) Prove that the perpendicular bisectors of a triangle meet at one point.
- (2) Prove the tangent-secant form $PT^2 = PA \cdot PB$.
- (3) In a cyclic quadrilateral, prove that opposite angles sum to 180° .
- (4) Let ABC be an equilateral triangle of side length s , and let P be any point in its closed interior. Use a 60° rotation and the triangle inequality to prove

$$PA + PB + PC \geq \sqrt{3}s,$$

with equality at the center of the triangle.

- (5) Let ABC be a triangle and D lie on BC . Prove Stewart's theorem.
- (6) Find a synthetic proof and a coordinate proof of the fact that medians of a triangle meet in a ratio $2 : 1$.
- (7) (**Putnam 1998 A1**) A right circular cone has base radius 1 and height 3. A cube is inscribed with one face in the base of the cone. Find the side length of the cube.

§14 Coordinate Geometry, Vectors, and Analytic Methods

Coordinates are not a last resort. They are a language for turning incidence, distance, angle, and area into algebra. Vectors are even more flexible because they keep geometry visible while allowing computation. The Putnam often rewards a coordinate setup chosen with symmetry: place a point at the origin, put a line on an axis, or normalize a circle to radius 1.

Definition 14.1 Dot product

For vectors $u = (u_1, \dots, u_n)$ and $v = (v_1, \dots, v_n)$, the dot product is

$$u \cdot v = u_1 v_1 + \dots + u_n v_n.$$

It satisfies $u \cdot v = \|u\| \|v\| \cos \theta$.

The dot product converts angle questions into algebra. Perpendicularity becomes $u \cdot v = 0$, projection becomes a scalar multiple, and length becomes $u \cdot u$.

Example 14.1. Median length formula


Let ABC be a triangle with side lengths $a = BC$, $b = CA$, $c = AB$. If m_a is the median from A to BC , prove that

$$m_a^2 = \frac{2b^2 + 2c^2 - a^2}{4}.$$

Solution. Put the midpoint of BC at the origin and write $B = v$, $C = -v$, and $A = u$. Then $a = 2\|v\|$, while $c^2 = \|u - v\|^2$ and $b^2 = \|u + v\|^2$. Adding gives

$$b^2 + c^2 = \|u + v\|^2 + \|u - v\|^2 = 2\|u\|^2 + 2\|v\|^2.$$

Since $m_a = \|u\|$ and $a^2 = 4\|v\|^2$, rearrangement gives the formula.



Theorem 14.1 Cauchy-Schwarz as vector geometry

For vectors $u, v \in \mathbb{R}^n$,

$$|u \cdot v| \leq \|u\| \|v\|.$$

Proof. If $v = 0$ the result is clear. Otherwise the vector $u - tv$ has nonnegative squared length for every real t . The quadratic $\|u - tv\|^2$ therefore has nonpositive discriminant, which gives $(u \cdot v)^2 \leq \|u\|^2 \|v\|^2$.

Example 14.2. Distance to a line

Find the distance from (x_0, y_0) to the line $ax + by + c = 0$.

Solution. The vector (a, b) is normal to the line. Choose any point Q on the line. The signed distance from $P = (x_0, y_0)$ to the line is the projection of $P - Q$ onto the unit normal $(a, b)/\sqrt{a^2 + b^2}$. Since $aQ_x + bQ_y + c = 0$,

this projection is

$$\frac{ax_0 + by_0 + c}{\sqrt{a^2 + b^2}}.$$

Taking absolute value gives the distance.

Example 14.3. Choosing coordinates wisely

Prove that the diagonals of a rectangle have equal length.

Solution. Place the rectangle at $(0,0)$, $(a,0)$, (a,b) , and $(0,b)$. The two diagonals have squared lengths $a^2 + b^2$ and $a^2 + b^2$. They are equal. The coordinate placement uses the symmetry of the rectangle and removes irrelevant translation and rotation.

Affine transformations

An affine transformation preserves lines, parallelism, ratios along a line, and midpoints, but it does not preserve lengths or angles. Therefore it is useful for incidence and ratio problems, especially with triangles and parallelograms. It is not appropriate when circles or perpendicularity are essential.

Example 14.4. Centroid by vectors

Let A, B, C have position vectors a, b, c . Prove that the medians meet at the point with position vector $(a + b + c)/3$.

Solution. The midpoint of BC has vector $(b + c)/2$. The point

$$g = \frac{a + b + c}{3}$$

lies on the line from A to this midpoint because

$$g = \frac{1}{3}a + \frac{2}{3} \cdot \frac{b + c}{2}.$$

The same expression is symmetric in a, b, c , so g lies on all three medians. It divides each median in the ratio $2 : 1$ from the vertex.

Example 14.5. A determinant area formula

For points (x_1, y_1) , (x_2, y_2) , and (x_3, y_3) , the signed area of the triangle is

$$\frac{1}{2} \det \begin{pmatrix} x_1 & y_1 & 1 \\ x_2 & y_2 & 1 \\ x_3 & y_3 & 1 \end{pmatrix}.$$

The determinant vanishes exactly when the three points are collinear. This formula is often the quickest way to turn geometry into algebra.

Example 14.6. Projection as an optimization

Find the point on the line through the origin in direction v closest to a point u .

Solution. Points on the line have form tv . Minimize

$$\|u - tv\|^2 = \|u\|^2 - 2t(u \cdot v) + t^2\|v\|^2.$$

The quadratic is minimized at $t = (u \cdot v)/\|v\|^2$. Thus the closest point is the projection of u onto the line. This calculation is a model for many least-squares and distance problems.

Problem patterns for this section

Analytic methods are strongest when the coordinate choice removes variables. Put a circle at the origin with radius 1. Put a line on an axis. Put the midpoint of a segment at the origin. Use vectors when ratios and midpoints matter; use dot products when angles and lengths matter; use determinants when area and collinearity matter. The setup is half the solution.

Proof-writing details for this section

Coordinate proofs should explain why the coordinate choice loses no generality. After setting coordinates, translate the geometric condition into equations and keep track of which variables are free. A clean coordinate solution usually ends by interpreting the algebra back in geometric language.

Exercises for Section 14

- (1) Prove the parallelogram identity $\|u + v\|^2 + \|u - v\|^2 = 2\|u\|^2 + 2\|v\|^2$.
- (2) Find the equation of the circle through $(0, 0)$, $(1, 0)$, and $(0, 1)$.
- (3) Use vectors to prove that the diagonals of a parallelogram bisect each other.
- (4) Derive the formula for the area of a triangle from a determinant.
- (5) Let P vary on the unit circle. Maximize the distance from P to a fixed point A .
- (6) Prove that four distinct points in the plane cannot have all six pairwise distances equal.
- (7) (**Putnam 2006 A1**) Find the volume of the region of points (x, y, z) satisfying $(x^2 + y^2 + z^2 + 8)^2 \leq 36(x^2 + y^2)$.

§15 Linear Algebra Essentials for Problem Solving

Linear algebra turns many problems about equations, transformations, and dimensions into statements about kernels and images. The Putnam often uses linear algebra in disguise: a polynomial interpolation problem is a dimension problem, a recurrence is a matrix power problem, and a combinatorial incidence relation is a vector space problem.

Definition 15.1 Rank and nullity

For a linear map $T : V \rightarrow W$, the rank is $\dim \operatorname{im} T$ and the nullity is $\dim \ker T$.

Theorem 15.1 Rank-nullity theorem

If V is finite-dimensional, then

$$\dim V = \operatorname{rank} T + \dim \ker T.$$

Proof. Choose a basis $v_1, \dots, v_k$ of $\ker T$ and extend it to a basis $v_1, \dots, v_k, u_1, \dots, u_r$ of V . The vectors $T(u_1), \dots, T(u_r)$ form a basis of the image: they span because every vector in V is a linear combination of the chosen basis, and they are independent because any linear relation among their images would put the corresponding linear combination of the u_i in the kernel, contradicting the basis extension. Thus $\operatorname{rank} T = r$ and $\dim V = k + r$.

Example 15.1. Polynomial interpolation as dimension

Show that there is a unique polynomial of degree at most n passing through any prescribed values at $n + 1$ distinct real numbers.

Solution. Let V be the vector space of polynomials of degree at most n , so $\dim V = n + 1$. Define $T : V \rightarrow \mathbb{R}^{n+1}$ by

$$T(P) = (P(x_0), P(x_1), \dots, P(x_n)),$$

where the x_i are distinct. If $T(P) = 0$, then P has $n + 1$ distinct roots, so $P = 0$. Thus T has trivial kernel. By rank-nullity, T has rank $n + 1$, so it maps onto $\mathbb{R}^{n+1}$. Existence and uniqueness follow.

Theorem 15.2 Trace and eigenvalues

For a square matrix over $\mathbb{C}$, the trace equals the sum of the eigenvalues counted with algebraic multiplicity, and the determinant equals their product.

This theorem is often used as a shortcut. If a problem gives $A^2 = A$, then eigenvalues must satisfy $\lambda^2 = \lambda$, so they are 0 or 1. The trace then equals the rank.

Example 15.2. Idempotent matrices

Let A be a real $n \times n$ matrix satisfying $A^2 = A$. Prove that $\operatorname{tr}(A)$ is a nonnegative integer.

Solution. The polynomial $x^2 - x = x(x - 1)$ has distinct roots, so A is diagonalizable over $\mathbb{R}$ with eigenvalues only 0 and 1. The trace is the sum of the eigenvalues, hence the number of eigenvalues equal to 1. Therefore $\operatorname{tr}(A)$ is a nonnegative integer. In fact it equals $\operatorname{rank}(A)$.

Example 15.3. A determinant trick

Compute the determinant of the $n \times n$ matrix with a on the diagonal and b off the diagonal.

Solution. Let J be the all-ones matrix. The given matrix is $(a - b)I + bJ$. The vector $(1, \dots, 1)$ is an eigenvector with eigenvalue $a - b + nb = a + (n - 1)b$. Every vector whose coordinates sum to zero is killed by J , so it is an eigenvector with eigenvalue $a - b$. Therefore the determinant is

$$(a - b)^{n-1}(a + (n - 1)b).$$

Example 15.4. Dimension forces a polynomial relation

Let V be the vector space of polynomials of degree at most n . Show that any $n + 2$ polynomials in V are linearly dependent.

Solution. The space V has basis $1, x, x^2, \dots, x^n$, so its dimension is $n + 1$. Any collection of more than $n + 1$ vectors in an $(n + 1)$ -dimensional vector space is linearly dependent. This simple dimension count is often hidden inside interpolation problems.

Working over finite fields

Linear algebra works over any field, including $\mathbb{F}_2$. Over $\mathbb{F}_2$, addition means xor, so toggling problems, parity configurations, and subsets are often vector space problems. A move set spans the reachable configurations. An invariant is a linear functional that vanishes on all moves.

Example 15.5. A toggling invariant

Three lamps are in a row. A move toggles either lamps 1 and 2 together or lamps 2 and 3 together. Which configurations are reachable from all off?

Solution. Work in $\mathbb{F}_2^3$. The two moves are vectors $(1, 1, 0)$ and $(0, 1, 1)$. Their span contains

$$(0, 0, 0), (1, 1, 0), (0, 1, 1), (1, 0, 1).$$

These are exactly the configurations with an even number of lamps on. The invariant is the parity of the total number of on lamps, since each move toggles two lamps.

Theorem 15.3 Cayley-Hamilton theorem

Every square matrix satisfies its own characteristic polynomial.

For Putnam work, Cayley-Hamilton is most useful in small dimensions. A 2×2 matrix A with trace t and determinant d satisfies

$$A^2 - tA + dI = 0.$$

This can reduce high powers of A to linear combinations of A and I .

Example 15.6. Reducing powers of a matrix

If A is a 2×2 matrix with trace 3 and determinant 2, then $A^2 = 3A - 2I$. Hence every higher power of A can be reduced recursively. For example,

$$A^3 = A(3A - 2I) = 3A^2 - 2A = 3(3A - 2I) - 2A = 7A - 6I.$$

This is the matrix analogue of reducing powers modulo a polynomial.

Problem patterns for this section

Linear algebra appears whenever there are many equations of the same type. If unknowns are coefficients, use dimension. If a process toggles states, use a vector space over $\mathbb{F}_2$. If a matrix satisfies a polynomial equation, inspect eigenvalues or use Cayley-Hamilton. If a trace appears, look for eigenvalues or diagonal entries counted in two ways. Rank-nullity turns impossibility into dimension comparison.

Proof-writing details for this section

Linear algebra solutions should name the vector space, the linear map, and the basis or projection being used. When rank, trace, or determinant appears, explain why it is invariant under the chosen change of basis. For finite-field encodings, state explicitly which operations are taking place modulo 2 or modulo p .

Exercises for Section 15

- (1) Prove that a square matrix with a zero row has determinant 0.
- (2) Find the eigenvalues of the all-ones matrix J_n .
- (3) Let $A^2 = 0$. Prove that $\text{tr}(A) = 0$.
- (4) Prove that if $AB = I$ for square matrices, then $BA = I$.
- (5) Let P be a projection matrix, meaning $P^2 = P$. Prove that $\text{rank}(P) = \text{tr}(P)$ without using diagonalization.
- (6) Let $v_1, \dots, v_m \in \mathbb{R}^n$ with $m > n$. Prove that there is a nontrivial linear relation among them. Deduce that for any distinct real numbers $t_1, \dots, t_{n+2}$ there are real coefficients $c_1, \dots, c_{n+2}$, not all zero, such that

$$\sum_{i=1}^{n+2} c_i t_i^j = 0 \quad \text{for every } j = 0, 1, \dots, n.$$

- (7) (**Putnam 2006 B4**) Let Z be the set of all 0-1 vectors in $\mathbb{R}^n$. If V is a k -dimensional linear subspace of $\mathbb{R}^n$, prove that $|V \cap Z| \leq 2^k$, and give an example attaining equality.

§16 Sequences, Series, and Asymptotic Thinking

Analysis on the Putnam often begins with estimates. Exact formulas are rare; useful upper and lower bounds are common. When a sequence is defined recursively, ask whether it is monotone and bounded. When a series appears, ask whether it telescopes, compares to a known series, or can be grouped dyadically.

Definition 16.1 Asymptotic comparison

We write $a_n = O(b_n)$ if there is a constant C such that $|a_n| \leq C|b_n|$ for all sufficiently large n . We write $a_n \sim b_n$ if $a_n/b_n \rightarrow 1$.

Asymptotics are not just notation. They prevent over-computation. A Putnam problem may only need to know that a quantity grows faster than another, not its exact value.

Theorem 16.1 Monotone convergence theorem for sequences

Every monotone increasing sequence that is bounded above converges. Every monotone decreasing sequence that is bounded below converges.

Proof. Let (a_n) be increasing and bounded above. The set of its values has a least upper bound L . For any $\varepsilon > 0$, the number $L - \varepsilon$ is not an upper bound, so some $a_N > L - \varepsilon$. Since the sequence is increasing, $L - \varepsilon < a_n \leq L$ for all $n \geq N$, proving convergence to L .

Example 16.1. A recursive limit

Let $a_1 = 1$ and $a_{n+1} = \sqrt{2 + a_n}$. Prove that a_n converges and find the limit.

Solution. First $1 \leq a_n \leq 2$ by induction: if $a_n \leq 2$, then $a_{n+1} \leq \sqrt{4} = 2$. Also $a_{n+1} \geq a_n$ is equivalent to $2 + a_n \geq a_n^2$, or $(a_n - 2)(a_n + 1) \leq 0$, true for $1 \leq a_n \leq 2$. Thus the sequence is increasing and bounded, so it converges. If the limit is L , then $L = \sqrt{2 + L}$, so $L^2 - L - 2 = 0$. Since $L \geq 0$, we get $L = 2$.

Example 16.2. A telescoping estimate

Prove that

$$\sum_{n=1}^{\infty} \frac{1}{n(n+1)} = 1.$$

Solution. Use partial fractions:

$$\frac{1}{n(n+1)} = \frac{1}{n} - \frac{1}{n+1}.$$

The partial sums telescope:

$$\sum_{n=1}^N \frac{1}{n(n+1)} = 1 - \frac{1}{N+1}.$$

Letting $N \rightarrow \infty$ gives the result.

Theorem 16.2 Integral comparison

If f is positive and decreasing on $[1, \infty)$, then

$$\int_1^{N+1} f(x) dx \leq \sum_{n=1}^N f(n) \leq f(1) + \int_1^N f(x) dx.$$

This theorem gives the convergence of $\sum 1/n^p$ for $p > 1$ and divergence for $p \leq 1$. More importantly, it gives the right order of growth for partial sums.

Example 16.3. Dyadic grouping

Prove that the harmonic series diverges by grouping terms.

Solution. Group terms as

$$1 + \frac{1}{2} + \left(\frac{1}{3} + \frac{1}{4}\right) + \left(\frac{1}{5} + \cdots + \frac{1}{8}\right) + \cdots.$$

In the block from $2^k + 1$ to 2^{k+1} , there are 2^k terms, each at least $1/2^{k+1}$, so the block sum is at least $1/2$. Infinitely many blocks contribute at least $1/2$, so the series diverges.

Theorem 16.3 Ratio test

For a series $\sum a_n$ with positive terms, if $\lim a_{n+1}/a_n = L < 1$, then the series converges. If $L > 1$, it diverges.

Proof. If $L < 1$, choose r with $L < r < 1$. Eventually $a_{n+1} \leq ra_n$, so the tail is bounded by a geometric series. If $L > 1$, then eventually $a_{n+1} \geq a_n$ by a factor greater than 1 often enough that the terms do not tend to zero, so the series cannot converge.

Example 16.4. Estimating a binomial coefficient

Show that

$$\binom{2n}{n} \geq \frac{4^n}{2n+1}.$$

Solution. The sum of all binomial coefficients in row $2n$ is 4^n . There are $2n+1$ coefficients, and the largest one is the central coefficient $\binom{2n}{n}$. Hence

$$4^n \leq (2n+1) \binom{2n}{n}.$$

This is an averaging estimate. It is crude but often sufficient.

Asymptotic proof style

When estimating, state the range clearly. Constants do not matter in O notation, but dependencies do. If a constant depends on a fixed parameter, say so. If it depends on n , it is not a constant for asymptotic purposes.

Example 16.5. A squeeze limit

If $0 \leq a_n \leq 1/n$ for all n , then $a_n \rightarrow 0$. This is the squeeze theorem. It is simple, but it is often the final move after a difficult estimate reduces an expression to a known vanishing bound.

Problem patterns for this section

For sequences, first test monotonicity and boundedness. If exact limits are hard, estimate above and below. For series, look for telescoping before applying tests. If terms resemble an integral, compare to an integral. If terms come in blocks, group them. In asymptotic problems, identify the dominant contribution and prove that the rest is smaller by a clear bound.

Proof-writing details for this section

Sequence proofs need both a guess and a control mechanism. If a limit is guessed from a differential equation or dominant term, the proof must show boundedness, monotonicity, Cauchy behavior, or an error estimate. For asymptotics, write which lower-order terms are being ignored and why they are negligible.

Exercises for Section 16

- (1) Determine whether $\sum 1/n^2$ converges using integral comparison.
- (2) Prove that the harmonic series diverges.
- (3) Let $a_{n+1} = (a_n + 2/a_n)/2$ with $a_1 > 0$. Guess the limit and prove convergence.
- (4) Show that $\sum_{n=1}^N 1/n = O(\log N)$.
- (5) Prove that if $a_n \geq 0$ and $\sum a_n$ converges, then $a_n \rightarrow 0$.
- (6) Let (a_n) be a sequence satisfying $0 < a_n < 1$ and

$$a_{n+1} \leq a_n - a_n^2$$

for every n . Prove that $a_n \rightarrow 0$, and then prove the quantitative estimate $a_n \leq 1/(n - 1 + 1/a_1)$.

- (7) **(Putnam 2006 B6)** Let $k > 1$ be an integer, $a_0 > 0$, and $a_{n+1} = a_n + 1/\sqrt[k]{a_n}$. Evaluate $\lim_{n \rightarrow \infty} a_n^{k+1}/n^k$.

§17 Real Analysis Methods in Problem Solving

Real analysis gives compactness, continuity, and approximation. These are existence tools. If a continuous function on a compact interval has values of opposite sign, it has a zero. If a continuous function on a compact set is real-valued, it attains a maximum and minimum. If a sequence is bounded, it often has a convergent subsequence. Putnam analysis problems frequently ask you to combine these facts with algebraic insight.

Theorem 17.1 Intermediate value theorem

If f is continuous on $[a, b]$ and y lies between $f(a)$ and $f(b)$, then there is $c \in [a, b]$ with $f(c) = y$.

Example 17.1. A fixed point on an interval

Let $f : [0, 1] \rightarrow [0, 1]$ be continuous. Prove that f has a fixed point.

Solution. Consider $g(x) = f(x) - x$. Since $f(0) \geq 0$, we have $g(0) \geq 0$. Since $f(1) \leq 1$, we have $g(1) \leq 0$. By the intermediate value theorem, there is $c \in [0, 1]$ such that $g(c) = 0$, so $f(c) = c$.

Theorem 17.2 Extreme value theorem

A continuous real-valued function on a compact interval $[a, b]$ attains its maximum and minimum.

The theorem is often used silently. If a problem asks you to minimize a continuous expression over a closed bounded set, you are allowed to select a minimizing point. Once selected, that point may satisfy symmetry, derivative, or perturbation conditions.

Example 17.2. A maximum with a perturbation

Let $x_1, \dots, x_n \geq 0$ and $x_1 + \dots + x_n = 1$. Prove that $x_1 x_2 \cdots x_n \leq n^{-n}$.

Solution. The product is continuous on the compact simplex, so it attains a maximum. If one coordinate is zero, the product is zero, not maximal unless all products are zero, which is false for equal positive coordinates. Thus at a maximum all coordinates are positive. If two coordinates x_i, x_j are unequal, replacing them by their average keeps the sum fixed and increases their product because

$$\left(\frac{x_i + x_j}{2}\right)^2 - x_i x_j = \frac{(x_i - x_j)^2}{4} > 0.$$

This increases the total product, a contradiction. Hence all coordinates are equal to $1/n$, giving maximum n^{-n} .

Definition 17.1 Compactness heuristic

When a problem has a bounded closed domain and a continuous quantity, choose an extremal point. Then perturb it. If the domain is not compact, first prove that extremal behavior cannot escape to infinity or to a missing boundary.

Example 17.3. Density of rational approximations

Prove that between any two real numbers $a < b$ there is a rational number.

Solution. Choose a positive integer n so large that $1/n < b - a$. There is an integer m with $na < m < nb$, for instance by taking the least integer greater than na . Then $a < m/n < b$, and m/n is rational. This simple approximation tool is often enough when a problem wants a rational or integer object close to a real one.

Theorem 17.3 Bolzano-Weierstrass theorem

Every bounded sequence of real numbers has a convergent subsequence.

This theorem is a compactness tool. It allows you to pass from infinitely many approximate objects to one limiting object. In contest solutions, it often appears in a simpler nested-interval form: repeatedly divide a bounded interval and keep a subinterval containing infinitely many terms.

Example 17.4. A subsequence argument

Let (a_n) be a bounded sequence with exactly one subsequential limit L . Prove that $a_n \rightarrow L$.

Solution. If a_n did not converge to L , then for some $\varepsilon > 0$ infinitely many terms would satisfy $|a_n - L| \geq \varepsilon$. Those terms form a bounded subsequence, which has a convergent subsequence by Bolzano-Weierstrass. Its limit would be a subsequential limit different from L , contradiction.

Theorem 17.4 Uniform continuity on compact intervals

If f is continuous on $[a, b]$, then for every $\varepsilon > 0$ there is $\delta > 0$ such that $|x - y| < \delta$ implies $|f(x) - f(y)| < \varepsilon$ for all $x, y \in [a, b]$.

Uniform continuity lets one choose a single δ that works everywhere. This matters when many intervals, partitions, or approximating points are involved.

Example 17.5. Equal values separated by a fixed distance

Let $f : [0, 1] \rightarrow \mathbb{R}$ be continuous with $f(0) = f(1)$. Prove that for each positive integer n , there exists $x \in [0, 1 - 1/n]$ with $f(x) = f(x + 1/n)$.

Solution. Define $g(x) = f(x + 1/n) - f(x)$ for $x \in [0, 1 - 1/n]$. If some $g(x) = 0$, we are done. Otherwise all values of g have the same sign by continuity. But summing at the points $0, 1/n, \dots, (n-1)/n$ gives

$$\sum_{k=0}^{n-1} \left(f\left(\frac{k+1}{n}\right) - f\left(\frac{k}{n}\right) \right) = f(1) - f(0) = 0.$$

The summands cannot all be strictly positive or all strictly negative. Hence one is zero or the continuous function changes sign, giving the desired x .

Approximation with rationals and integers

If a real problem has an integer conclusion, approximate real quantities by rationals only after deciding which inequalities have room for error. A strict inequality can survive small perturbations; a non-strict equality may not. This is why epsilon margins appear in rigorous approximation arguments.

Problem patterns for this section

Continuity problems often reduce to a sign change, an attained extremum, or a limiting subsequence. Compactness is the tool that lets you choose a best object. If the domain is not compact, prove that the relevant behavior stays in a compact region. When a statement asks for equal values at shifted inputs, sum differences around a cycle and use the intermediate value theorem.

Proof-writing details for this section

Real analysis proofs often hide in compactness and continuity. Whenever a maximum, minimum, or limiting point is used, state the theorem that guarantees it. If an interval is split into pieces, check endpoint behavior. The Putnam frequently rewards a short continuity argument, but only when all hypotheses are visible.

Exercises for Section 17

- (1) Prove that every odd-degree real polynomial has a real root.
- (2) Let $f : [0, 1] \rightarrow \mathbb{R}$ be continuous with $f(0) = f(1)$. Prove that for every positive integer n , there exists $x \in [0, 1 - 1/n]$ such that $f(x) = f(x + 1/n)$.
- (3) Prove that a continuous injective function on an interval is strictly monotone.
- (4) Let $f : [0, 1] \rightarrow [0, 1]$ be continuous and satisfy $f(f(x)) = x$ for every x . Prove that either $f(x) = x$ for every x , or f is strictly decreasing.
- (5) Show that a continuous function on $\mathbb{R}$ with finite limits at $\pm\infty$ is bounded.
- (6) Let f be continuous on $[0, 1]$ and differentiable on $(0, 1)$ with $f(0) = f(1) = 0$. Prove that for every real a , the function $g_a(x) = f(x) - ax(1 - x)$ has a critical point in $(0, 1)$.
- (7) **(Putnam 2004 A1)** A basketball player has season totals $S(N)$ made shots in her first N attempts. Early on $S(N) < 0.8N$, later $S(N) > 0.8N$. Must there have been an N with $S(N) = 0.8N$?

§18 Calculus and Functional Methods

Calculus is a problem-solving language for change. Derivatives prove monotonicity, convexity, and local optimality. Integrals average functions and convert sums into areas. Functional methods study equations where the unknown is a function. The Putnam often combines these: a functional equation may need continuity, differentiability, or monotonicity extracted from the equation itself.

Theorem 18.1 Mean value theorem

If f is continuous on $[a, b]$ and differentiable on (a, b) , then there exists $c \in (a, b)$ such that

$$f'(c) = \frac{f(b) - f(a)}{b - a}.$$

Proof. Apply Rolle's theorem to the function obtained by subtracting the secant line from f . The new function has equal values at a and b , so it has a critical point. At that point, the derivative of f equals the slope of the secant line.

Example 18.1. A monotonicity proof

Prove that $\sin x < x$ for all $x > 0$.

Solution. For $x > 0$, apply the mean value theorem to $\sin t$ on $[0, x]$. There exists $c \in (0, x)$ such that

$$\frac{\sin x - 0}{x} = \cos c.$$

Since $\cos c < 1$ for $c \in (0, x)$ unless $c = 0$, we get $\sin x/x < 1$ for $0 < x$ in a range where c is not a multiple of 2π . More generally, the function $g(x) = x - \sin x$ has derivative $1 - \cos x \geq 0$ and $g(0) = 0$, so $g(x) > 0$ for $x > 0$.

Theorem 18.2 Convexity test

If $f''(x) \geq 0$ on an interval, then f is convex there. If $f''(x) \leq 0$, then f is concave.

Example 18.2. A logarithmic inequality

Prove that for $x > -1$,

$$\ln(1+x) \leq x.$$

Solution. Define $g(x) = x - \ln(1+x)$. Then $g'(x) = 1 - 1/(1+x) = x/(1+x)$. The derivative is negative on $(-1, 0)$ and positive on $(0, \infty)$, so g has a global minimum at 0. Since $g(0) = 0$, we get $g(x) \geq 0$, which is the desired inequality.

Definition 18.1 Functional equation strategy

When solving for a function, first plug in special values such as 0, 1, or fixed points. Then compare the equation at two inputs. If regularity is given, use it late; first extract algebraic consequences.

Example 18.3. Cauchy's equation with continuity

Find all continuous functions $f : \mathbb{R} \rightarrow \mathbb{R}$ satisfying

$$f(x+y) = f(x) + f(y)$$

for all real x, y .

Solution. Setting $y = 0$ gives $f(0) = 0$. Repeated addition gives $f(nx) = nf(x)$ for positive integers n , and using $0 = f(x + (-x))$ gives $f(-x) = -f(x)$. Thus $f(q) = qf(1)$ for every rational q . For real x , choose rational numbers $q_n \rightarrow x$. Continuity gives $f(x) = \lim f(q_n) = \lim q_n f(1) = xf(1)$. Hence $f(x) = cx$ for some constant c , and every such function works.

Theorem 18.3 Rolle's theorem

If f is continuous on $[a, b]$, differentiable on (a, b) , and $f(a) = f(b)$, then some $c \in (a, b)$ satisfies $f'(c) = 0$.

Proof. By the extreme value theorem, f attains a maximum and a minimum on $[a, b]$. If both occur only at endpoints, then the equal endpoint values force the function to be constant at its extremes, and any interior point works if the function is constant. Otherwise an interior maximum or minimum exists, and Fermat's theorem gives derivative zero there.

Example 18.4. Repeated roots force derivative roots

If a polynomial P has k distinct real roots, prove that P' has at least $k - 1$ distinct real roots.

Solution. Order the roots $r_1 < \cdots < r_k$. On each interval $[r_i, r_{i+1}]$, the polynomial has equal endpoint values, namely zero. Rolle's theorem gives a point $c_i \in (r_i, r_{i+1})$ with $P'(c_i) = 0$. These $k - 1$ points are distinct.

Example 18.5. Integral average

If f is continuous on $[a, b]$, prove that there is $c \in [a, b]$ such that

$$f(c) = \frac{1}{b-a} \int_a^b f(x) dx.$$

Solution. The extreme value theorem gives $m \leq f(x) \leq M$ on $[a, b]$. Integrating gives

$$m \leq \frac{1}{b-a} \int_a^b f(x) dx \leq M.$$

By the intermediate value theorem, f attains every value between m and M , so it attains its average value.

Functional equations with regularity

Regularity assumptions such as continuity, boundedness on an interval, monotonicity, or differentiability usually eliminate wild solutions. First solve the equation on integers and rationals. Then use the regularity to pass from rationals to reals. This pattern appears in Cauchy's equation and many of its variants.

Example 18.6. Differentiating a functional equation

Suppose f is differentiable and $f(x+y) = f(x)f(y)$ for all real x, y , with $f(0) = 1$ and f positive. Setting $y = 0$ is consistent. Differentiating with respect to y at $y = 0$ gives $f'(x) = f(x)f'(0)$. Thus f satisfies a differential equation, so $f(x) = e^{cx}$ where $c = f'(0)$.

Problem patterns for this section

Calculus problems often require only one derivative, but it must be the right derivative. To prove an inequality, move everything to one side and study monotonicity or convexity. To relate roots and critical points, use Rolle's theorem. To solve a differentiable functional equation, first plug in special values, then differentiate with respect to one variable after justifying that the equation holds in a differentiable setting.

Proof-writing details for this section

Calculus proofs should not skip the sign analysis. After differentiating, identify the interval on which the derivative has a sign, and connect that sign to monotonicity or convexity. For functional equations, first determine special values such as $f(0)$, then use regularity assumptions such as continuity or differentiability to rule out wild solutions.

Exercises for Section 18

- (1) Use derivatives to prove $e^x \geq 1 + x$.
- (2) Find the maximum of $x(1 - x)$ on $[0, 1]$.
- (3) Prove that $\arctan x \leq x$ for $x \geq 0$.
- (4) Find all differentiable functions satisfying $f'(x) = f(x)$ and $f(0) = 1$.
- (5) Solve all continuous $f : \mathbb{R} \rightarrow \mathbb{R}$ such that $f(x + y) = f(x)f(y)$ and $f(0) = 1$, with f positive.
- (6) Let f be differentiable and satisfy $f(x + y) = f(x) + f(y) + xy$. Determine all such functions.
- (7) **(Putnam 2002 B3)** Prove that for every integer $n > 1$, $\frac{1}{2ne} < \frac{1}{e} - (1 - 1/n)^n < \frac{1}{ne}$.

§19 Polynomials, Complex Numbers, and Algebraic Methods

Polynomials are rigid. A polynomial of degree n cannot have too many roots, cannot oscillate arbitrarily, and is determined by enough values. Complex numbers add geometry to algebra: multiplication rotates and scales, roots of unity encode symmetry, and conjugation reflects.

Theorem 19.1 Remainder theorem

For a polynomial $P(x)$, the remainder upon division by $x - a$ is $P(a)$. In particular, $x - a$ divides $P(x)$ if and only if $P(a) = 0$.

Proof. Polynomial division gives $P(x) = (x - a)Q(x) + r$ where r is constant. Substituting $x = a$ gives $P(a) = r$.

Example 19.1. A polynomial divisibility test

Find all real a such that $x^2 + ax + 1$ divides $x^4 + x^2 + 1$.

Solution. If r is a root of $x^2 + ax + 1$, then $r^2 = -ar - 1$. Reduce powers using this relation. We have

$$r^4 + r^2 + 1 = (r^2)^2 + r^2 + 1 = (-ar - 1)^2 - ar - 1 + 1.$$

A cleaner route is factorization:

$$x^4 + x^2 + 1 = (x^2 + x + 1)(x^2 - x + 1).$$


Thus $x^2 + ax + 1$ must be one of these factors, so $a = 1$ or $a = -1$.

Definition 19.1 Roots of unity

The complex solutions of $z^n = 1$ are

$$1, \omega, \omega^2, \dots, \omega^{n-1}, \quad \omega = e^{2\pi i/n}.$$

They form a regular n -gon on the unit circle.



Example 19.2. A roots of unity filter

Let S be the sum of coefficients of x^{3k} in $(1 + x)^n$. Express S using roots of unity.

Solution. Let $\omega = e^{2\pi i/3}$. The filter

$$\frac{1 + \omega^j + \omega^{2j}}{3}$$

is 1 if $3 \mid j$ and 0 otherwise. Therefore

$$S = \frac{(1+1)^n + (1+\omega)^n + (1+\omega^2)^n}{3}.$$

This method turns a divisibility condition on exponents into an average over roots of unity.

Theorem 19.2 Fundamental theorem of algebra

Every nonconstant polynomial with complex coefficients has at least one complex root. Hence a degree n polynomial has exactly n complex roots counted with multiplicity.

For Putnam purposes, the main consequence is that factoring over $\mathbb{C}$ is always possible. This lets us compare roots, apply conjugation, and convert coefficient conditions into symmetric sums.

Theorem 19.3 Vieta's formulas

If

$$P(x) = x^n + c_{n-1}x^{n-1} + \cdots + c_0$$

has roots $r_1, \dots, r_n$ in $\mathbb{C}$, counted with multiplicity, then

$$r_1 + \cdots + r_n = -c_{n-1}, \quad r_1 r_2 \cdots r_n = (-1)^n c_0,$$

and the intermediate coefficients are signed elementary symmetric sums of the roots.

Vieta's formulas are the bridge between coefficients and roots. They are especially useful when a problem gives integer coefficients, because algebraic roots then satisfy strong symmetric restrictions.

Example 19.3. Root location from coefficients

Let $P(x) = x^2 - 5x + 6$. Without factoring first, find the sum and product of its roots.

Solution. Vieta gives sum 5 and product 6. The roots are therefore 2 and 3, but the key information was available before finding them. In harder problems, the roots may not be expressible simply, while their sum and product are still useful.

Example 19.4. Lagrange interpolation idea

Let P have degree at most 2 and values $P(0) = 1$, $P(1) = 2$, $P(2) = 5$. Find $P(3)$.

Solution. The first differences are 1 and 3, so the second difference is 2. A quadratic with constant second difference 2 has next first difference 5, so $P(3) = 10$. Equivalently, interpolation gives $P(x) = x^2 + 1$.

Complex-number geometry

In the complex plane, multiplication by $e^{i\theta}$ rotates by angle θ , multiplication by a real scalar scales, and conjugation reflects across the real axis. This makes complex numbers natural for rotations, regular polygons, and circle equations.

Example 19.5. Rotation by complex multiplication

If z is a complex number, then iz is the point obtained by rotating z by 90° counterclockwise about the origin. Thus a square with one vertex at 0 and another at z has possible third vertices $z + iz$ and $z - iz$. This gives coordinate-free algebra for square and equilateral-triangle constructions.

Example 19.6. A root-counting trap

If a nonzero polynomial of degree at most n has $n + 1$ roots, it must be the zero polynomial. This statement is so common that it is easy to misuse: the roots must be distinct, and the object must really be a polynomial of bounded degree. Rational functions, absolute values, and piecewise expressions need separate checks.

Problem patterns for this section

Polynomial problems have three main languages: coefficients, roots, and values. Vieta moves between coefficients and roots. Interpolation moves between values and coefficients. The remainder theorem moves between divisibility and values. Complex numbers add a geometric language for rotation and symmetry. When a problem gives too many conditions, count degrees. When it gives symmetry, use roots of unity or conjugation.

Proof-writing details for this section

Polynomial and complex-number arguments depend on exact structural facts: degree, roots, multiplicity, conjugation, and leading coefficient. When using roots of unity or interpolation, write the filtering or uniqueness statement explicitly. When using logarithmic derivatives, separate the case where the point is already a root.

Exercises for Section 19

- (1) Prove that if P has real coefficients and z is a complex root, then $\bar{z}$ is also a root.
- (2) Find all roots of $z^4 = 1$ and draw them.
- (3) Use the remainder theorem to find the remainder when x^{100} is divided by $x^2 - 1$.
- (4) Prove that $x^n - y^n$ is divisible by $x - y$.
- (5) Let P be a polynomial of degree n with $P(k) = 0$ for $k = 1, \dots, n$. Express $P(n+1)$ in terms of its leading coefficient.
- (6) Use roots of unity to find the sum of coefficients of terms whose exponents are congruent to 1 modulo 3 in $(1+x)^n$.
- (7) (**Putnam 2005 A3**) Let $p(z)$ be a polynomial of degree n whose zeros all have absolute value 1, and put $g(z) = p(z)/z^{n/2}$. Prove that every zero of $g'(z)$ has absolute value 1.

§20 Functional Equations, Invariants, and Strategy Synthesis

The final Putnam skill is synthesis. A problem may begin as a functional equation, then require modular arithmetic, then finish with an invariant. The best solvers keep several languages available and switch when the current language stops producing information.

Functional equations are particularly good training for synthesis because they reward experimentation. Plug in special values, compare two substitutions, iterate the equation, and look for fixed points. If a function maps integers to integers, use divisibility and descent. If it maps reals to reals and has regularity, use continuity or monotonicity after extracting algebra.

Example 20.1. A function on integers with a hidden invariant

Find all functions $f : \mathbb{Z} \rightarrow \mathbb{Z}$ such that

$$f(x + y) = f(x) + f(y) + 2xy$$

for all integers x, y .

Solution. The term $2xy$ suggests completing a square. Define $g(x) = f(x) - x^2$. Then

$$g(x + y) = f(x + y) - (x + y)^2 = f(x) + f(y) + 2xy - x^2 - 2xy - y^2 = g(x) + g(y).$$

Thus g is additive on $\mathbb{Z}$, so $g(x) = cx$ where $c = g(1)$ is an integer. Therefore

$$f(x) = x^2 + cx.$$

Every such function works. The invariant was not a quantity preserved by a process; it was a transformation that preserved additivity.

Synthesis checklist

- Extract simple consequences before using advanced tools.
- Check small cases and equality cases.
- Encode relations as graphs, polynomials, or matrices when pairwise data appears.
- Look for an invariant under the operation in the problem.
- If the problem has a maximum or minimum hidden inside, choose an extremal object.
- When stuck, solve a weaker version and inspect what information was missing.

Example 20.2. Synthesis with subsets

Let $a_1, \dots, a_n$ be integers. Prove that there is a nonempty subset $I \subseteq \{1, \dots, n\}$ such that $\sum_{i \in I} a_i$ is divisible by n , provided the a_i are allowed to be reordered and consecutive blocks are considered after the reordering.

Solution. In fact no reordering is needed if we allow consecutive blocks. Consider partial sums $s_k = a_1 + \dots + a_k$ modulo n . If some s_k is 0, we are done. If not, two partial sums have the same residue, and their difference is a consecutive block sum divisible by n . A consecutive block is a subset, so the result follows. Here the synthesis is small but important: subset divisibility was solved by sequence partial sums and pigeonhole.

Example 20.3. An invariant in a coloring process

A row of n lamps is initially off. A move chooses a consecutive block and toggles every lamp in that block. Which configurations are reachable?

Solution. Every configuration is reachable. Toggle the block from lamp 1 through lamp k whenever the desired final state changes between positions k and $k + 1$, and also toggle a prefix if the first lamp must be on. More explicitly, represent a configuration by a binary vector $(b_1, \dots, b_n)$. The interval vectors span $\mathbb{F}_2^n$ because the vector with a single 1 in position k is the sum modulo 2 of the prefix intervals $[1, k]$ and $[1, k - 1]$. Thus every binary vector is reachable. This is a process problem solved by linear algebra over $\mathbb{F}_2$.

Example 20.4. A mixed divisibility and polynomial problem

Let $P(x)$ be a polynomial with integer coefficients. Prove that if $a \equiv b \pmod{m}$, then $P(a) \equiv P(b) \pmod{m}$.

Solution. It is enough to prove the result for monomials. Since $a \equiv b \pmod{m}$, repeated multiplication gives $a^k \equiv b^k \pmod{m}$ for every $k \geq 0$. Multiplying by the integer coefficient of each monomial and summing preserves congruence. This simple fact combines polynomial structure with modular arithmetic and is often the hidden reason that a polynomial modulo m is well-defined on residue classes.

Example 20.5. A full synthesis problem

Let S be a finite nonempty set of integers. Suppose that for every $a, b \in S$, the integer $a + b$ is even and $(a + b)/2 \in S$. Prove that S has one element.

Solution. Choose two elements $u, v \in S$ with maximal distance $|u - v|$. If $u \neq v$, then their average $w = (u + v)/2$ belongs to S and lies strictly between u and v . This does not contradict maximality directly, since w is closer to both endpoints. Instead choose two distinct elements with minimal positive distance. Their average is an integer in S and is strictly between them, giving a smaller positive distance, contradiction. Thus no two distinct elements exist. The problem combines parity, closure, finiteness, and an extremal choice.

Switching languages deliberately

A good solution may say: “We now view this as a graph,” or “We now work modulo p ,” or “Consider the polynomial whose roots are the selected numbers.” These sentences are not cosmetic. They tell the reader why a new tool is relevant. On the Putnam, a clear language switch can make a solution much easier to grade.

Example 20.6. A graph-linear algebra synthesis

Let G be a graph with vertices $1, \dots, n$. A move chooses a vertex and toggles the colors of all vertices adjacent to it. Over $\mathbb{F}_2$, the move vectors are the columns of the adjacency matrix of G . A target coloring is reachable from the all-off coloring exactly when its vector lies in the column space of this matrix. Invariants are vectors in the left nullspace. This single translation converts a game into rank-nullity.

How to finish a hard problem

When you have partial progress, identify the exact missing implication. Do not keep redoing computations. Write “It remains to prove ...” and make that sentence as narrow as possible. Many final ideas appear only after the missing claim is isolated.

Problem patterns for this section

Mixed problems are solved by reduction chains. Write each reduction as a clean sentence: the original problem is equivalent to a congruence, the congruence is controlled by a polynomial identity, the identity

follows from roots, and the equality case is handled by an extremal choice. A chain is only as strong as its weakest link, so make sure every transition is reversible or clearly one-directional as needed.

Proof-writing details for this section

Synthesis problems require a proof spine. Write the invariant, the algebraic relation, and the extremal choice as separate sentences before combining them. This keeps the solution from becoming a maze. If the same object is viewed in two languages, for example as a graph and as a vector space, state the dictionary between the languages.

Exercises for Section 20

- (1) Solve all functions $f : \mathbb{Z} \rightarrow \mathbb{Z}$ such that $f(x + y) = f(x) + f(y)$.
- (2) Find all functions $f : \mathbb{R} \rightarrow \mathbb{R}$ satisfying $f(x + y) = f(x) + f(y)$ and f is bounded on some interval.
- (3) A finite graph has each vertex colored red or blue. A move changes the color of one vertex and all its neighbors. Interpret reachability as a linear algebra problem over $\mathbb{F}_2$.
- (4) Let S be a finite nonempty set of nonnegative integers such that $|a - b| \in S$ whenever $a, b \in S$. Prove that either $S = \{0\}$ or there are positive integers m and q such that

$$S = \{0, m, 2m, \dots, qm\}.$$

- (5) Let $f : \mathbb{N} \rightarrow \mathbb{N}$ be bounded and satisfy $f(n + f(n)) = f(n)$ for every n . If the positive integer M satisfies $f(n) \leq M$ for all n , prove that $f(n + L) = f(n)$ for all n , where $L = \text{lcm}(1, 2, \dots, M)$. Thus f is periodic.
- (6) Create a Putnam-style problem that combines a polynomial identity, a congruence, and an extremal choice. Then solve it.
- (7) (**Putnam 2001 A1**) A binary operation $*$ on a set S satisfies $(x * y) * x = y$ for all $x, y \in S$. Prove that $x * (y * x) = y$ for all $x, y \in S$.

§A Classical Theorems Every Putnam Student Should Know

A theorem is useful on the Putnam only if you know when to try it. This appendix is a field guide. The goal is not to memorize a large list but to build recognition. Each theorem below is paired with triggers, warnings, and a miniature use case.

Number theory toolkit

Euclidean algorithm. Try it when a problem involves gcds, coprime variables, or linear combinations. The key fact is that $\gcd(a, b) = \gcd(b, a - bq)$. It enables Bezout coefficients and modular inverses.

Fermat and Euler. Try them when exponents occur modulo an integer and the base is coprime to the modulus. Beware that Euler's theorem gives a period dividing $\varphi(n)$, not always the minimal period.

Chinese remainder theorem. Try it when congruence conditions have coprime moduli. It lets you solve locally and assemble globally.

Valuations. Try them when divisibility by powers of primes is the issue. They convert multiplication into addition and lcm or gcd questions into max or min of exponents.

Example A.1. Using Bezout cleanly

If $\gcd(a, b) = 1$, prove that every integer solution of $ax + by = 0$ has the form $x = bt$, $y = -at$.

Solution. From $ax = -by$, we get $b \mid ax$. Since $\gcd(a, b) = 1$, Euclid's lemma gives $b \mid x$. Write $x = bt$. Then $abt + by = 0$, so $y = -at$.

Combinatorics toolkit

Pigeonhole. Try it when many objects have few possible labels, residues, sizes, or signatures.

Double counting. Try it when a sum of local quantities might equal a global quantity.

Inclusion-exclusion. Try it when avoiding forbidden properties or counting objects with at least one property.

Generating functions. Try them when objects are built by independent choices whose sizes add.

Algebra and analysis toolkit

Polynomial identity principle. Try it when a polynomial has many roots or many prescribed values.

Cauchy-Schwarz. Try it when sums of squares, reciprocal sums, or dot products appear.

AM-GM. Try it when positive variables have a fixed sum or product.

Compactness and continuity. Try them when an extremum or a zero should exist but cannot be written down explicitly.

Mean value theorem. Try it when a difference quotient, monotonicity, or a derivative bound appears.

A practical theorem sheet

A useful theorem sheet has three columns: theorem, trigger, and warning. For Cauchy-Schwarz, a trigger is a square of a sum or a reciprocal sum; a warning is that equality requires proportionality. For the Chinese remainder theorem, a trigger is simultaneous congruences; a warning is that the moduli should be coprime unless a compatibility condition is checked. This format trains recognition, not recitation.

Example A.2. One theorem used three ways

Cauchy-Schwarz can prove a geometric statement, an algebraic inequality, and an averaging bound. Geometrically, $|u \cdot v| \leq \|u\| \|v\|$. Algebraically, $(\sum a_i b_i)^2 \leq (\sum a_i^2)(\sum b_i^2)$. In Engel form, $\sum x_i^2/a_i \geq (\sum x_i)^2/(\sum a_i)$ for positive a_i . The theorem is the same, but the recognition pattern changes.

Theorems that create objects

Some theorems do not estimate or simplify; they create objects. The intermediate value theorem creates a root. The extreme value theorem creates a maximum or minimum. Hall's theorem creates a matching. The Chinese remainder theorem creates an integer with prescribed local behavior. When an existence problem has no obvious construction, look for a theorem that creates the object indirectly.

Theorems that destroy objects

Other theorems prove impossibility. Modular arithmetic can destroy integer solutions. Planar graph bounds destroy drawings of K_5 or $K_{3,3}$. A degree bound destroys a polynomial with too many roots. Compactness can destroy an infinite avoidance pattern by forcing a limit. Knowing whether a theorem creates or destroys helps choose the proof direction.

Tool	Best trigger	Common warning
AM-GM	Positive variables with fixed sum or product	Equality requires the averaged quantities equal.
Cauchy	Squares, reciprocals, dot products	Check proportionality for equality.
Jensen	Sum of one function at many inputs	Convexity must hold on the interval.
CRT	Several congruences	Non-coprime moduli need compatibility.
Euler/Fermat	High powers modulo n	Base must be invertible modulo n .
IVT	Continuous sign change	Domain must be an interval.
Rank-nullity	Too many vectors or equations	Must be over a field.

Example A.3. A theorem chain

Prove that for every positive integer n , there is a positive integer made only of digits 0 and 1 that is divisible by n .

Solution. Consider the n integers

$$1, \quad 11, \quad 111, \quad \dots, \quad \underbrace{11 \dots 1}_{n \text{ digits}}.$$

If any is divisible by n , we are done. Otherwise two have the same residue modulo n by pigeonhole. Their difference is divisible by n and has decimal form consisting of some 1s followed by some 0s. This number uses only digits 0 and 1. The proof uses a theorem chain: residues create boxes, pigeonhole creates two equal residues, subtraction creates the desired number.

How much theorem proof to know

You do not need to reprove every classical theorem during the contest, but you should know the proof idea. Knowing the proof idea prevents misuse. Euler's theorem is a permutation argument, so it needs invertibility. Jensen is a convexity argument, so the interval and equality conditions matter. Rank-nullity

is a basis-extension argument, so it needs finite-dimensional vector spaces. The proof idea tells you the hypotheses.

Exercises for Section A

- (1) For each theorem above, write one original problem where it is the main tool.
- (2) Find three problems in which Cauchy-Schwarz can be used in different forms.
- (3) Prove Euclid's lemma using Bezout's identity.
- (4) Prove Fermat's little theorem using the permutation idea from Euler's theorem.
- (5) Give an example where applying AM-GM too early loses the equality case.
- (6) Build a one-page personal theorem sheet with triggers rather than statements only.

§B Common Mistakes in Putnam Solutions

A Putnam solution is graded for proof, not for intention. The most common mistakes are not computational. They are logical gaps, hidden assumptions, and unjustified generalizations from examples. This appendix lists errors that strong students make and gives ways to repair them.

Mistake 1: assuming the extremum exists

If the domain is finite, an extremum exists. If the domain is compact and the function is continuous, an extremum exists. Otherwise you must prove existence or use an approximate extremum.

For example, a positive continuous function on $(0, 1)$ need not attain a minimum. The function $f(x) = x$ has infimum 0 but no minimum on $(0, 1)$. If a proof says “choose the smallest value” on a noncompact set, it may be invalid.

Mistake 2: dividing by something that might be zero

Every division creates a condition. If you divide by $x - y$, handle the case $x = y$ first. If you divide by a polynomial, prove it is nonzero under the hypotheses.

Example B.1. A dangerous cancellation

Solve $x^2 = x$.

Bad solution. Divide by x to get $x = 1$.

Repair. Rewrite as $x(x - 1) = 0$, so $x = 0$ or $x = 1$. The division lost the solution $x = 0$.

Mistake 3: proving only a necessary condition

A congruence obstruction can prove impossibility, but satisfying the congruence does not prove existence. A suspected answer must satisfy all necessary conditions and then be constructed or forced by a sufficient argument.

Mistake 4: changing the quantifiers

The statement “for every x there exists y ” is not the same as “there exists y for every x .” Many functional equation and combinatorics mistakes come from moving a chosen object outside the scope where it was chosen.

Mistake 5: using a theorem without checking hypotheses

Jensen needs convexity or concavity on the whole interval in use. Euler’s theorem needs coprimality. The mean value theorem needs continuity on the closed interval and differentiability inside. Equality conditions also have hypotheses.

Example B.2. A hidden assumption about order

A common false move is to assume variables are ordered without saying so. If an expression is symmetric, you may write “without loss of generality, assume $x \leq y \leq z$.” If it is not symmetric, that assumption may change the problem. The phrase “without loss of generality” is a claim that must be justified by symmetry or by a reversible relabeling.

Partial credit versus false claims

A correct lemma that does not finish the problem can earn credit. A false claim that seems to finish the problem can ruin the solution. If you are unsure, write the lemma you can prove and clearly state what remains. Do not hide uncertainty behind vague phrases like “clearly” or “it follows” when the implication is the main difficulty.

Example B.3. A quantifier repair

The statement “for every positive integer n there is a prime larger than n ” is true. The statement “there is a prime larger than every positive integer n ” is false. The same words appear, but the order of quantifiers changes everything. In solutions, keep chosen constants inside the paragraph where they are chosen unless you prove they work uniformly.

A final proof audit

Before submitting, check these items: all cases from divisions and square roots, all theorem hypotheses, all equality cases, all uses of finiteness or compactness, and the exact conclusion. The audit should be fast because it is mechanical. It is not a search for new ideas; it is a search for preventable gaps.

Example B.4. A false pattern from small cases

The numbers $n^2 + n + 41$ are prime for $n = 0, 1, \dots, 39$, but at $n = 40$ the value is $40^2 + 40 + 41 = 41^2$. Small cases suggest conjectures, not proofs. When a pattern looks strong, search for the structural reason it must continue. If no such reason appears, test the input that makes a factor visible.

Example B.5. A missing endpoint

Suppose a proof maximizes $f(x) = x(1 - x)$ on $(0, 1)$ by saying the maximum exists because the interval is bounded. This is invalid reasoning: boundedness of the interval is not enough; the domain is not closed. In this example a maximum does exist, but the proof has not justified it. On $[0, 1]$, the extreme value theorem applies immediately.

Danger words

Words such as “obvious,” “clearly,” and “without loss of generality” are not forbidden, but they are warning lights. After writing one of them, ask what precise fact makes the claim true. If the answer is symmetry, name the symmetry. If the answer is monotonicity, prove the monotonicity. If the answer is a standard theorem, check its hypotheses.

Example B.6. Equality case lost by squaring

If x is real and $\sqrt{x^2} = x$, then $x \geq 0$. Squaring both sides of an equation can introduce extraneous solutions, while taking square roots can lose signs. Whenever radicals appear, keep track of domains and sign restrictions. A clean solution often isolates a radical, squares, and then checks all candidates in the original equation.

Exercises for Section B

- (1) Find the flaw in this proof: since $a^2 = b^2$, divide by $a - b$ to get $a + b = 0$.

- (2) Give an example of a bounded sequence that does not converge.
- (3) Give an example of a continuous function on an open interval that does not attain a maximum.
- (4) Write a false statement created by switching “for all” and “there exists.”
- (5) Take one of your old solutions and mark every place where a theorem hypothesis is used.
- (6) Create a checklist for the last five minutes of writing a proof.

§C Timed Practice and Exam Strategy

The Putnam is a proof contest under time pressure. Strategy matters. The competition is not won by solving problems in printed order. It is won by recognizing which problems fit your strengths, extracting rigorous progress from them, and reserving enough time to turn ideas into readable proofs.

What changed in 2026

The mathematical work still occupies six hours, but it is now divided into four ninety-minute sessions with designated breaks. A student who exits the room during a session may not return to continue that session. Start times are being coordinated so that the sessions are synchronous across time zones. These rules make each ninety-minute endpoint a real deadline: unfinished reasoning cannot simply be carried on in the room after a brief unscheduled exit, and clean writing cannot always be postponed to a long second half.

The three-pass method for a ninety-minute session

First pass: orient. Read every problem in the current packet. Mark each as friendly, possible, or distant. Record one possible tool beside each problem, but do not let the survey become a collection of long starts.

Second pass: convert. Attack the friendliest problem. Once the proof spine is sound, begin the clean write-up while the details are fresh. A nearly complete proof is usually a better use of the session than a speculative start on a substantially harder problem.

Third pass: salvage and audit. Return to the best remaining lead. If a complete solution is not available, write rigorous partial progress: a proved lemma, reduction, construction, bound, or necessary condition. Reserve the final minutes to check hypotheses, exceptional cases, and the exact conclusion.

How to practice

Practice in four modes. In learning mode, spend unlimited time and read solutions after serious effort. In sprint mode, give yourself ten minutes to find a promising first move. In proof mode, write a complete solution to a problem you already understand. In session mode, work for exactly ninety minutes on a prepared packet, with only the materials allowed by current Putnam rules. Older released exams can still be used: divide each former three-hour half into two balanced practice packets, but treat that division as training rather than as a claim about the official 2026 packet order.

Example C.1. Triage example

Suppose a problem asks for a determinant, a second asks for a geometric maximum, and a third asks for a functional equation with no regularity. If linear algebra is your strength, the determinant is a friendly problem. The geometric maximum may become friendly after coordinates. The functional equation may be distant unless special substitutions quickly produce structure. This classification is personal, not absolute.

Writing advice. A solution should begin with the main idea as soon as possible. Avoid long preambles. Define notation before using it. Keep equations aligned. If a proof uses a lemma, state the lemma clearly and prove it. If you use a known theorem, name it and check its hypotheses in the sentence where it is used.

Point-saving habits

- Write complete sentences around key implications.
- Do not erase useful partial work unless it is wrong.

- State the final answer prominently only after proving it.
- If you cannot finish, state a correct reduced claim and prove it.
- Leave space between attempts so the grader can follow the final argument.

Building the new kind of stamina

The 2026 format requires repeated acceleration: focus for ninety minutes, stop on command, recover during a designated break, and restart on a new packet. Train both levels. A single-session rehearsal develops pacing. A full-day rehearsal of four ninety-minute sessions develops recovery, handwriting endurance, and the ability to reset after a disappointing packet. Once the official timetable is published, imitate its break lengths and start times rather than guessing them.

Example C.2. How to write a partial solution

If you cannot solve a problem asking for all possible n , write necessary conditions and constructions separately. A clear proof that every solution must satisfy $n \equiv 0 \pmod{3}$ is useful even if you cannot prove sufficiency. A construction for infinitely many n is useful even if you cannot prove they are all. Label the two directions so the grader can see the structure.

A weekly routine

A balanced week might include two technique drills, two single-problem deep dives, one proof rewrite, and one ninety-minute packet. Technique drills keep tools fresh. Deep dives build persistence. Proof rewrites improve communication. Timed packets train triage and conversion from idea to proof. Every few weeks, replace the single packet with two sessions separated by a planned break; closer to the competition, include at least one full four-session rehearsal.

Before entering each session

Use the designated break for practical needs, including the restroom, water or food as permitted, and a brief physical reset. Confirm with the supervisor what may enter the room. Because leaving mid-session ends your work for that session, do not treat an unscheduled exit as part of your pacing plan. Students who need an accommodation should arrange it through the official institutional process in advance.

During a ninety-minute session

If an approach produces no new mathematical information for roughly ten minutes, change language or change problems. A language change might mean trying a graph, a modulus, coordinates, a generating function, or an extremal object. The relevant test is progress, not mere activity: expanding expressions for ten minutes without obtaining a lemma, bound, pattern, or obstruction is a signal to switch.

Time	A flexible ninety-minute template
0 to 6 min	Read the current packet; mark friendly, possible, and distant.
6 to 42 min	Develop the strongest lead and begin its clean proof as soon as the spine is secure.
42 to 65 min	Finish that proof or work the next best problem if the first is complete.
65 to 80 min	Convert remaining progress into a rigorous solution or clearly labeled partial result.
80 to 90 min	Audit logic, hypotheses, edge cases, notation, and final conclusions.

This table is a starting point, not a command. A student who writes slowly should begin the clean proof earlier. A student with two short solutions should shorten the first work interval. The invariant is that proof-writing time is protected before the session ends.

The proof rewrite drill

Choose a problem you solved. Put the solution away. The next day, rewrite it from memory in under fifteen minutes. Then compare. Any step you cannot reconstruct is either too complicated or not understood. Rewrite again until the proof has a clear spine: setup, key lemma, main argument, conclusion.

Example C.3. Postmortem categories

After a timed block, classify each missed problem. A knowledge miss means a theorem or technique was unfamiliar. A recognition miss means the tool was known but not seen. An execution miss means the idea was found but the proof failed. A triage miss means too much time went to the wrong problem. These categories lead to different training responses.

Scoring mindset

A polished solution to one problem is usually worth more than fragments on many problems. However, correct fragments can matter when they are substantial and clearly connected to the problem. The balance is personal. During practice, track how often a partial attempt becomes a full solution after ten more minutes. This helps calibrate when to persist and when to move on.

Synchronous start times

Do not infer the 2026 local start time from an older Putnam schedule or from another time zone. Confirm the reporting time, building access, and permitted arrival window with the local supervisor. Plan transportation backward from that confirmed time. The point of synchronous scheduling is competition security, so an avoidable late arrival may not be repairable by starting later.

Exercises for Section C

- (1) Take a past problem and spend exactly ten minutes listing possible tools without solving it.
- (2) Rewrite a messy solution into a one-page clean proof.
- (3) For one week, do one problem per day in sprint mode and record the first useful observation.
- (4) Make a personal table of strengths: algebra, number theory, combinatorics, geometry, analysis, linear algebra.

- (5) Simulate one ninety-minute session and write a postmortem identifying time lost to unproductive paths and time lost by delaying the clean write-up.
- (6) Rehearse two ninety-minute sessions separated by a designated break. Record what you did during the break and whether the second session began with full attention.
- (7) Before the competition, complete one four-session rehearsal using the official timetable then in force. Afterward, separate mathematical errors from pacing, recovery, and logistics errors.

§D A Roadmap Through Past Putnam Problems

Past problems are best used as a map, not as a pile. Random practice can be discouraging because the Putnam spans many areas and difficulty levels. A better plan is to group problems by technique, then revisit the same technique at increasing difficulty. The official MAA Putnam Archive lists recent problems and solutions by year, and it is the best starting point for current official materials.

How to use the archive

Choose a theme from the list below. For each theme, select three to five problems from different years. Spend at least twenty minutes on each before reading a solution. After reading, write a technique tag, such as “partial sums modulo n ” or “polynomial roots.” The tag matters more than the year.

Theme	What to look for in past problems
Algebra	Factorizations, polynomial identities, symmetric expressions, interpolation.
Number theory	Congruences, valuations, modular orders, Diophantine descent.
Combinatorics	Double counting, pigeonhole, graph encoding, inclusion-exclusion.
Analysis	Continuity, compactness, inequalities, sequences, integrals.
Linear algebra	Rank, trace, determinant tricks, vector spaces over finite fields.
Geometry	Coordinate setups, power of a point, transformations, convexity.

A twelve-week route

(1) Weeks 1 and 2: proof forms, algebraic manipulation, and inequalities.

(2) Weeks 3 and 4: congruences, Diophantine equations, and arithmetic functions.

(3) Weeks 5 and 6: counting, pigeonhole, generating functions, and graph encoding.

(4) Weeks 7 and 8: geometry, vectors, and linear algebra.

(5) Weeks 9 and 10: sequences, real analysis, calculus, and functional equations.

(6) Weeks 11 and 12: mixed timed sets, solution writing, and error analysis.

Example D.1. Turning an archive problem into a lesson

If a past problem is solved by considering partial sums modulo n , do not file it only as “number theory.” File it as “pigeonhole plus modular residues.” Then look for a graph problem or sequence problem where the same partial-sum idea appears. This cross-tagging builds transfer, which is the real goal of Putnam preparation.

After each past problem

Write five lines: the first useful observation, the failed idea that seemed plausible, the key lemma, the theorem used, and the shortest memory tag. For example: “Sort variables; naive AM-GM fails; smoothing lemma; convexity; equality at boundary.”

Archive tags

Use tags such as modular obstruction, partial sums, Cauchy, compactness, Vieta jumping, roots of unity, rank-nullity, and extremal graph. A problem can and should have more than one tag. The second tag is often the one that teaches transfer.

Example D.2. A sample archive entry

Title: polynomial values problem. First observation: too many prescribed values suggest interpolation. Failed idea: expand the polynomial directly. Key lemma: a degree n polynomial with $n + 1$ roots is zero. Memory tag: roots determine polynomial. Follow-up: find another problem where prescribed integer values imply divisibility.

Spacing past problems

Do not solve ten problems on the same theme in one sitting and then abandon the theme. Solve three, wait a week, solve two more, then mix the theme with other areas. Spacing builds recall. Mixing builds recognition. Both are necessary because the exam does not announce which technique is intended.

Reading solutions productively

Read a solution only after you can name your obstruction. When reading, stop at the first idea you did not find. Close the solution and try to finish from there. This converts reading from passive consumption into guided problem solving.

Month	Main focus	Archive practice
Month 1	Algebra, proof, inequalities	A1/B1 style problems with short solutions.
Month 2	Number theory and combinatorics	Problems tagged residues, pigeonhole, double counting.
Month 3	Analysis, linear algebra, geometry	Mixed A2/B2/A3 problems with full writeups.
Final weeks	Simulation and review	Several ninety-minute packets and at least one four-session rehearsal.

Difficulty ladders

For each technique, build a ladder. A level 1 problem applies the theorem directly. A level 2 problem requires choosing the theorem. A level 3 problem combines the theorem with another area. A level 4 problem hides the theorem behind a transformation. When using the archive, do not jump only to the hardest level. A ladder builds speed and recognition.

Example D.3. A technique ladder for partial sums

Level 1: prove a consecutive block sum is divisible by n . Level 2: prove a subset sum is divisible by n by ordering or adding dummy terms. Level 3: combine partial sums with inequalities to control the size of the block. Level 4: use partial sums in a group, vector space, or polynomial coefficient setting. The same idea becomes more powerful as the setting changes.

Maintaining the archive notebook

Each entry should be short enough to review quickly. Record the problem label, topic, technique tags, the key sentence of the proof, and one similar problem to try later. Do not copy full official solutions into the notebook. Rewrite the idea in your own words. The act of compression is what makes the archive useful.

Exercises for Section D

- (1) Pick five past A1 or B1 problems and tag each by method rather than topic.
- (2) Pick three problems from different years that use polynomials. Identify the shared idea.
- (3) Pick three problems that use modular arithmetic. Record the modulus and why it was natural.
- (4) Build a personal archive of ten solved problems, each with a one-sentence memory tag.
- (5) Re-solve one old problem after two weeks without looking at your notes.
- (6) Adapt one released twelve-problem competition into four balanced ninety-minute practice packets. Keep a note that this is a training arrangement, then use the official 2026 packet order when it becomes available.

§E Detailed Solutions and Proof Expansions

The purpose of this section is to make the proof details explicit. The main chapters emphasize intuition and technique; the following solutions show how those techniques become complete written arguments. The Putnam-labeled solutions are written in the same style as the rest of the manuscript, with the official problem labels retained for practice organization.

Section 1: General Strategies

Section 1, Exercise 4

Consider the partial sums $s_k = a_1 + \cdots + a_k$ for $1 \leq k \leq n$. If one of these is divisible by n , then the corresponding initial segment is the desired nonempty subset. Otherwise the n partial sums occupy only the $n - 1$ nonzero residue classes modulo n . Two of them, say s_i and s_j with $i < j$, must have the same residue. Their difference $a_{i+1} + \cdots + a_j$ is then divisible by n , and this is a nonempty subset sum. The distinctness of the original integers is not needed for this argument; the real engine is the pigeonhole principle applied to partial sums.

Putnam 2001 B1

Assign to each square a sign ϵ_{ij} , equal to 1 for red and -1 for black. The condition that each row has the same number of red and black squares says $\sum_j \epsilon_{ij} = 0$ for every row i , and the corresponding column condition says $\sum_i \epsilon_{ij} = 0$ for every column j . The difference between the red sum and the black sum is

$$\sum_{i=1}^n \sum_{j=1}^n \epsilon_{ij}((i-1)n + j).$$

Split this into the row part $\sum_i (i-1)n \sum_j \epsilon_{ij}$ and the column part $\sum_j j \sum_i \epsilon_{ij}$. Both vanish by the row and column balance conditions. Therefore the signed difference is zero, so the red entries and black entries have equal sum. The invariant is not the number of red squares, which is already fixed, but the signed first moment of the coloring with respect to the row-by-row labels.

Section 2: Proof Techniques

Section 2, Exercise 6

Let $s_k = a_1 + \cdots + a_k$ for $1 \leq k \leq n$. If some $s_k \equiv 0 \pmod{n}$, then the problem is finished. If no partial sum is divisible by n , then the n residues $s_1, \dots, s_n$ lie in the $n - 1$ nonzero residue classes modulo n . Hence two of them are equal modulo n . If $s_i \equiv s_j \pmod{n}$ with $i < j$, then $s_j - s_i = a_{i+1} + \cdots + a_j$ is divisible by n . This proof is a model of a contradiction-free pigeonhole argument: either the target appears directly, or the absence of the target forces a collision that produces it indirectly.

Putnam 1999 A5

Let V be the finite-dimensional real vector space of polynomials of degree at most 1999. Define

$$N(p) = \int_{-1}^1 |p(x)| dx.$$

This is a norm on V : positivity and homogeneity are immediate, the triangle inequality follows from the triangle inequality for absolute value, and $N(p) = 0$ implies that the continuous function $|p|$ is zero on $[-1, 1]$, hence p has infinitely many roots and must be the zero polynomial.

Choose the coefficient norm $\|p\|_c = \max_i |c_i|$ for $p(x) = \sum c_i x^i$. On the compact set $\|p\|_c = 1$, the function $N(p)$ is continuous and strictly positive, so it has a positive minimum δ . Therefore $N(p) \geq \delta \|p\|_c$ for every p by scaling. Since $|p(0)| = |c_0| \leq \|p\|_c$, we get $|p(0)| \leq \delta^{-1} N(p)$. Thus the required constant exists, for instance $C = 1/\delta$. The proof is really a compactness proof about norms on a finite-dimensional space.

Section 3: Algebraic Manipulation

Section 3, Exercise 5

Assume $a + b + c = 0$. Then $c = -(a + b)$, so

$$a^4 + b^4 + c^4 = a^4 + b^4 + (a + b)^4 = 2a^4 + 2b^4 + 4a^3b + 6a^2b^2 + 4ab^3.$$

On the other hand,

$$2(a^2b^2 + b^2c^2 + c^2a^2) = 2a^2b^2 + 2b^2(a + b)^2 + 2a^2(a + b)^2,$$

which expands to the same expression. A cleaner symmetric proof is to use $a + b + c = 0$ to get $a^2 + b^2 + c^2 = -2(ab + bc + ca)$, square both sides, and compare terms. The key detail is that the condition $a + b + c = 0$ converts fourth powers into pairwise products because the elementary symmetric sum $ab + bc + ca$ controls the whole expression.

Putnam 1999 A2

If $p = 0$, there is nothing to prove. Otherwise the leading coefficient of p is positive. Every real root of p has even multiplicity, since an odd-multiplicity real root would force the polynomial to change sign near that root. Hence the real-root part of p is a square of a real polynomial. The remaining irreducible real factors are positive quadratic factors of the form $(x - a)^2 + b^2$ with $b > 0$.

Each such quadratic is a sum of two squares, namely $(x - a)^2 + b^2$. The product of two sums of two squares is again a sum of two squares by the identity

$$(u^2 + v^2)(r^2 + s^2) = (ur - vs)^2 + (us + vr)^2.$$

Multiplying all the quadratic factors together and then multiplying by the square coming from the real roots gives a sum of two squares, after absorbing the positive leading coefficient as $(\sqrt{c})^2$. Thus p is a finite sum of squares of real polynomials. The proof works because nonnegativity forces exactly the factorization pattern that the two-square identity can preserve.

Section 4: Inequalities I

Section 4, Exercise 5

By Engel form of Cauchy,

$$\sum_{cyc} \frac{x}{y+z} = \sum_{cyc} \frac{x^2}{x(y+z)} \geq \frac{(x+y+z)^2}{x(y+z) + y(z+x) + z(x+y)}.$$

Since $x + y + z = 1$, the denominator is $2(xy + yz + zx)$. Also $xy + yz + zx \leq (x + y + z)^2/3 = 1/3$. Therefore

$$\sum_{cyc} \frac{x}{y+z} \geq \frac{1}{2(xy + yz + zx)} \geq \frac{3}{2}.$$

Equality requires $x = y = z = 1/3$. The proof is detailed enough for a contest solution because it explains both applications: Cauchy creates the reciprocal lower bound, and the elementary symmetric inequality controls the denominator.

Putnam 2003 A2

If some $a_i + b_i = 0$, then both a_i and b_i are zero, and both sides are zero. Otherwise set $c_i = a_i/(a_i + b_i)$, so $1 - c_i = b_i/(a_i + b_i)$ with $0 \leq c_i \leq 1$. Dividing the desired inequality by $((a_1 + b_1) \cdots (a_n + b_n))^{1/n}$ reduces it to

$$(c_1 \cdots c_n)^{1/n} + ((1 - c_1) \cdots (1 - c_n))^{1/n} \leq 1.$$

By AM-GM, $(c_1 \cdots c_n)^{1/n} \leq \frac{1}{n} \sum c_i$ and $((1 - c_1) \cdots (1 - c_n))^{1/n} \leq \frac{1}{n} \sum (1 - c_i)$. Adding the two inequalities gives the result. Equality occurs when all the c_i are equal, which means the ratios $a_i : b_i$ are the same for all i .

Section 5: Inequalities II

Section 5, Exercise 6

Let the fixed product be P . Write $a_i = P^{1/n} x_i$, so $x_1 \cdots x_n = 1$. It is enough to prove $x_1 + \cdots + x_n \geq n$. Instead of naming AM-GM, consider the function $\phi(t) = e^t$, which is convex. Put $x_i = e^{u_i}$. The product condition becomes $u_1 + \cdots + u_n = 0$. Jensen's inequality gives

$$\frac{1}{n} \sum_{i=1}^n e^{u_i} \geq e^{(u_1 + \cdots + u_n)/n} = 1.$$

Thus $\sum x_i \geq n$, and therefore $\sum a_i \geq nP^{1/n}$. Equality occurs when all $u_i = 0$, equivalently when all $a_i = P^{1/n}$. This proof shows the geometry behind AM-GM: after taking logarithms, equal product becomes fixed average.

Putnam 2006 B5

For each $x \in [0, 1]$, complete the square in the integrand:

$$x^2 f(x) - x(f(x))^2 = -x \left(f(x) - \frac{x}{2} \right)^2 + \frac{x^3}{4}.$$

The square term is nonpositive after multiplication by $-x$, so the expression is at most $x^3/4$ pointwise. Integrating gives

$$\int_0^1 x^2 f(x) dx - \int_0^1 x(f(x))^2 dx \leq \int_0^1 \frac{x^3}{4} dx = \frac{1}{16}.$$

The bound is attained by the continuous function $f(x) = x/2$. Hence the maximum is $1/16$. The pointwise square completion is stronger than a global inequality because it identifies the optimizing function at every value of x simultaneously.

Section 6: Divisibility and Congruences

Section 6, Exercise 6

Write $a = 1 + pt$. Then

$$1 + a + \cdots + a^{p-1} = \sum_{j=0}^{p-1} (1 + pt)^j.$$

Modulo p^2 , we may use $(1 + pt)^j \equiv 1 + jpt$, because every term of degree at least two contains p^2 . Hence

$$\sum_{j=0}^{p-1} (1 + pt)^j \equiv p + pt \sum_{j=0}^{p-1} j = p + pt \frac{p(p-1)}{2} \equiv p \pmod{p^2}.$$

Thus the sum is divisible by p but not by p^2 . The oddness of p ensures that the factor $(p-1)/2$ is an integer, and the displayed congruence shows exactly where the second power of p fails to appear.

Putnam 2000 B2

Let $g = \gcd(m, n)$. Using $\binom{n}{m} = \frac{n}{m} \binom{n-1}{m-1}$, the expression becomes

$$\frac{g}{m} \binom{n-1}{m-1}.$$

It remains to prove that m/g divides $\binom{n-1}{m-1}$. Since $\binom{n}{m}$ is an integer, for every prime p we have

$$v_p \left(\binom{n}{m} \right) = v_p(n) - v_p(m) + v_p \left(\binom{n-1}{m-1} \right) \geq 0.$$

If $v_p(m) > v_p(n)$, this gives $v_p \left(\binom{n-1}{m-1} \right) \geq v_p(m) - v_p(n)$; if $v_p(m) \leq v_p(n)$, there is no contribution from m/g . These inequalities are exactly the statement that m/g divides $\binom{n-1}{m-1}$. Therefore the original expression is an integer.

Section 7: Diophantine Equations

Section 7, Exercise 3

The equation $1/x + 1/y = 1/6$ becomes $6x + 6y = xy$ after multiplying by $6xy$. Rearranging gives

$$xy - 6x - 6y = 0, \quad (x-6)(y-6) = 36.$$

Thus positive integer solutions correspond to positive divisors d of 36 by setting $x-6 = d$ and $y-6 = 36/d$. Therefore

$$(x, y) = (6 + d, 6 + 36/d)$$

for each positive divisor d of 36. This divisor transformation is a standard Diophantine move: it turns a rational equation into a finite factor search.

Putnam 2001 A5

Suppose $a^{n+1} - (a+1)^n = 2001$. Reducing modulo a gives $-(a+1)^n \equiv -1 \equiv 2001 \pmod{a}$, so a divides $2002 = 2 \cdot 7 \cdot 11 \cdot 13$. Since $2001 \equiv 0 \pmod{3}$, the residue of a modulo 3 is restricted. If $a \equiv 0 \pmod{3}$, then $(a+1)^n$ is not divisible by 3; if $a \equiv 2 \pmod{3}$, then $a+1$ is divisible by 3 but a^{n+1} is not. Both are impossible. Hence $a \equiv 1 \pmod{3}$, and then $(a+1)^n \equiv 2^n \pmod{3}$ must equal 1, so n is even.

If a were even, then modulo 4 the left side would be $-(a+1)^n \equiv -1 \pmod{4}$, since n is even and $a+1$ is odd. But $2001 \equiv 1 \pmod{4}$, contradiction. Thus a is odd. Since a divides 2002, it divides $1001 = 7 \cdot 11 \cdot 13$. The congruence $a \equiv 1 \pmod{3}$ rules out the factor 11, because $11 \equiv 2 \pmod{3}$, so a divides $7 \cdot 13$. Also reducing modulo 4 with n even gives $a^{n+1} - (a+1)^n \equiv a \pmod{4}$, so $a \equiv 1 \pmod{4}$. Among the divisors of $7 \cdot 13$, this leaves only $a = 1$ and $a = 13$.

The value $a = 1$ gives $1 - 2^n$, never 2001. For $a = 13$, the equation is $13^{n+1} - 14^n = 2001$. The parity argument already gives n even. When $n > 2$ is even, $13^{n+1} \equiv 5 \pmod{8}$ and $14^n \equiv 0 \pmod{8}$, so the left side is $5 \pmod{8}$, while $2001 \equiv 1 \pmod{8}$. Hence n cannot exceed 2. Finally $13^3 - 14^2 = 2197 - 196 = 2001$, so the unique solution is $(a, n) = (13, 2)$.

Section 8: Arithmetic Functions

Section 8, Exercise 2

If $n = \prod p_i^{e_i}$, then the number of positive divisors is $d(n) = \prod (e_i + 1)$. This product is odd exactly when every factor $e_i + 1$ is odd, which is exactly when every exponent e_i is even. But all exponents in the prime factorization are even exactly when n is a perfect square. The proof is a good example of why multiplicative functions are most transparent after prime factorization.

Putnam 2003 B3

It is enough to compare the exponent of each prime p on both sides. The exponent of p in $n!$ is Legendre's sum

$$\sum_{e \geq 1} \left\lfloor \frac{n}{p^e} \right\rfloor.$$

Now examine the right side. The exponent of p in $\text{lcm}\{1, 2, \dots, \lfloor n/i \rfloor\}$ is the largest e such that $p^e \leq n/i$. Equivalently, it is the number of integers $e \geq 1$ for which $i \leq n/p^e$. Therefore the total exponent of p in the product over i is

$$\sum_{i=1}^n \sum_{e \geq 1} \mathbf{1}_{i \leq n/p^e} = \sum_{e \geq 1} \left\lfloor \frac{n}{p^e} \right\rfloor.$$

This matches the exponent in $n!$ for every prime p , so the two positive integers are equal. The proof is a double count of pairs (i, e) , where p^e appears in the lcm attached to i .

Section 9: Counting Methods

Section 9, Exercise 4

There are 3^n functions from an n -element set to a 3-element set. To count the onto functions, subtract those that miss at least one target value. For a fixed target value, 2^n functions miss it, and there are three choices for the missed value. If two target values are missed, then the function is forced to land entirely in the remaining value, giving 1 function for each of the three pairs of missed values. Inclusion-exclusion gives

$$3^n - 3 \cdot 2^n + 3.$$

The last term is added back because constant functions were subtracted twice when we excluded functions missing a specified value.

Putnam 2006 A4

For each position k , let I_k be the indicator of the event that the permutation has a local maximum at k . The desired average is $\mathbb{E}(I_1 + \cdots + I_n)$, which equals $\sum \mathbb{E}(I_k)$ by linearity of expectation. Independence is not required.

At an endpoint, the entry is a local maximum precisely when it is larger than its only neighbor, which has probability $1/2$. At an interior position, the three values in positions $k-1, k, k+1$ are equally likely to appear in any relative order, and the middle position is a local maximum precisely when its value is the largest of the three. This has probability $1/3$. Therefore the average number of local maxima is

$$\frac{1}{2} + \frac{1}{2} + (n-2)\frac{1}{3} = \frac{n+1}{3}.$$

The proof is short because expectation lets us count local events without understanding their dependence.

Section 10: Pigeonhole and Extremal Principles**Section 10, Exercise 5**

Write each chosen integer uniquely as $2^k r$, where r is odd. The odd part r lies between 1 and $2n$ and is odd, so there are exactly n possible odd parts. Choosing $n+1$ integers forces two of them to have the same odd part, say $2^a r$ and $2^b r$. If $a \leq b$, then $2^a r$ divides $2^b r$. This is a classic pigeonhole proof in which the boxes are not intervals but chains under divisibility.

Putnam 2002 A2

Choose two of the five points and draw a plane through them and the center of the sphere. This plane cuts the sphere into two closed hemispheres, with the chosen two points lying on the boundary of both. The remaining three points lie in the two closed half-spheres determined by the plane. By the pigeonhole principle, at least two of these three points lie on the same side, allowing boundary points to be counted in either hemisphere. Those two points together with the original two boundary points lie in one closed hemisphere. Hence some four of the five points lie on a closed hemisphere.

Section 11: Recurrences and Generating Functions**Section 11, Exercise 5**

The characteristic equation of $a_{n+2} = 3a_{n+1} - 2a_n$ is $r^2 - 3r + 2 = 0$, so $r = 1$ or $r = 2$. Thus $a_n = A + B2^n$. The initial conditions give $A + B = 0$ and $A + 2B = 1$, hence $B = 1$ and $A = -1$. Therefore $a_n = 2^n - 1$. Checking the recurrence directly confirms the formula: $2^{n+2} - 1 = 3(2^{n+1} - 1) - 2(2^n - 1)$.

Putnam 1999 A3

The generating function gives $a_0 = 1$, $a_1 = 2$, and $a_{n+2} = 2a_{n+1} + a_n$. Let

$$M = \begin{pmatrix} 0 & 1 \\ 1 & 2 \end{pmatrix}, \quad v = \begin{pmatrix} 1 \\ 2 \end{pmatrix}.$$

Then $M^n v = (a_n, a_{n+1})^T$. Since M is symmetric, the squared length of this vector is

$$a_n^2 + a_{n+1}^2 = (M^n v)^T (M^n v) = v^T M^{2n} v.$$

But $M^{2n}v = (a_{2n}, a_{2n+1})^T$, so

$$v^T M^{2n}v = a_{2n} + 2a_{2n+1} = a_{2n+2},$$

using the recurrence once more. Therefore $a_n^2 + a_{n+1}^2 = a_{2n+2}$, so the required integer is $m = 2n + 2$. The matrix method makes the hidden norm identity visible.

Section 12: Graph Theory

Section 12, Exercise 6

We prove by induction on the number of vertices that every tournament has a directed Hamiltonian path. For one vertex this is clear. Suppose a tournament on $n - 1$ vertices has a directed path $v_1 \rightarrow v_2 \rightarrow \cdots \rightarrow v_{n-1}$. Insert the remaining vertex w into this path. If $w \rightarrow v_1$, put w at the beginning. If $v_{n-1} \rightarrow w$, put w at the end. Otherwise, there is a first index i such that $w \rightarrow v_i$. Since w does not beat v_1 , this index is greater than 1, and by minimality $v_{i-1} \rightarrow w$. Then

$$v_1 \rightarrow \cdots \rightarrow v_{i-1} \rightarrow w \rightarrow v_i \rightarrow \cdots \rightarrow v_{n-1}$$

works. The proof is an insertion algorithm, not just an existence proof.

Putnam 2012 B3

Make a bipartite graph whose left vertices are the $2n - 1$ days and whose right vertices are the $2n$ teams. Join a day to every team that won its game on that day. We want a matching that covers all days, so Hall's theorem says it is enough to prove that any set D of d days has at least d neighboring winning teams. Let W be the set of teams that won at least once on the days in D , and write $w = |W|$. Suppose for contradiction that $w < d$. A team outside W lost every game it played on those d days. Hence two teams outside W could not have played each other on any of those days, because one of them would have won and entered W . Therefore every game played by a team outside W on a day in D was against a team in W . There are $2n - w$ teams outside W , and each plays once on each of the d days, giving $d(2n - w)$ cross games between W and its complement during those days. But in a round-robin tournament, a fixed outside team can play a fixed team in W at most once, so the total number of such cross games is at most $w(2n - w)$. Since $2n - w > 0$, this gives $d \leq w$, contradicting $w < d$. Hall's condition holds, so the desired choice of winning teams exists.

Section 13: Euclidean Geometry

Section 13, Exercise 6

Place the triangle at vectors A, B, C . The midpoint of BC is $(B + C)/2$, so the median from A consists of points $A + t((B + C)/2 - A)$. The centroid candidate is $G = (A + B + C)/3$. It lies on this median because

$$G = A + \frac{2}{3} \left(\frac{B + C}{2} - A \right).$$

Thus $AG : G((B + C)/2) = 2 : 1$. The same calculation holds for the other medians, so all three medians meet at G and are divided in the ratio $2 : 1$. The vector proof is concise because averaging the three vertices gives the common point directly.

Putnam 1998 A1

Take a vertical cross-section through the cone's axis and through a diagonal of the cube's bottom face. The cone becomes an isosceles triangle of height 3 and base 2, while the cube becomes a rectangle of height s and width $s\sqrt{2}$, where s is the cube side length. At height s above the base, the cone's horizontal radius has shrunk linearly from 1 to $1 - s/3$.

The top face of the cube has diagonal $s\sqrt{2}$, so half of that diagonal must equal the available radius at height s . Therefore

$$\frac{s\sqrt{2}}{2} = 1 - \frac{s}{3}.$$

Solving gives

$$s \left(\frac{\sqrt{2}}{2} + \frac{1}{3} \right) = 1, \quad s = \frac{6}{3\sqrt{2} + 2} = \frac{3\sqrt{2}}{3 + \sqrt{2}}.$$

The whole three-dimensional problem reduces to one similar-triangles relation in the correct cross-section.

Section 14: Coordinates and Vectors

Section 14, Exercise 6

Suppose four distinct points were pairwise at distance $d > 0$. Translate one point to the origin, and let $u, v, w \in \mathbb{R}^2$ be the position vectors of the other three. Then

$$\|u\| = \|v\| = \|w\| = d.$$

The condition $\|u - v\| = d$ gives

$$d^2 = \|u - v\|^2 = \|u\|^2 + \|v\|^2 - 2u \cdot v = 2d^2 - 2u \cdot v,$$

so $u \cdot v = d^2/2$. The same calculation gives

$$(1) \quad u \cdot w = v \cdot w = \frac{d^2}{2}.$$

The vectors u and v are linearly independent, because their angle has cosine $(u \cdot v)/d^2 = 1/2$ and is therefore 60° . Write $w = \alpha u + \beta v$. Taking dot products with u and v and using (1) gives

$$\alpha + \frac{\beta}{2} = \frac{1}{2}, \quad \frac{\alpha}{2} + \beta = \frac{1}{2}.$$

Thus $\alpha = \beta = 1/3$. Consequently

$$\|w\|^2 = \frac{1}{9}\|u + v\|^2 = \frac{1}{9}(\|u\|^2 + \|v\|^2 + 2u \cdot v) = \frac{d^2}{3},$$

contradicting $\|w\|^2 = d^2$. Hence no four distinct pairwise equidistant points exist in the plane.

Putnam 2006 A1

Use cylindrical coordinates with $\rho^2 = x^2 + y^2$. The inequality becomes

$$(\rho^2 + z^2 + 8)^2 \leq 36\rho^2.$$

Both sides are nonnegative, and the left square root is positive, so this is equivalent to

$$\rho^2 + z^2 + 8 \leq 6\rho.$$

Rearranging gives

$$(\rho - 3)^2 + z^2 \leq 1.$$

Thus the region is obtained by rotating the disk of radius 1 centered at $(\rho, z) = (3, 0)$ around the z -axis. By Pappus's theorem, or by direct cylindrical integration, the volume is the area π of the disk times the distance $2\pi \cdot 3$ traveled by its center. The volume is therefore $6\pi^2$.

Section 15: Linear Algebra

Section 15, Exercise 5

If $P^2 = P$, then every vector v decomposes as $v = Pv + (v - Pv)$. The first term lies in $\text{im } P$, and the second lies in $\ker P$ because $P(v - Pv) = Pv - P^2v = 0$. Also the intersection of $\text{im } P$ and $\ker P$ is zero: if $w = Pu$ and $Pw = 0$, then $w = P^2u = Pu = w$, so $w = 0$.

Choose a basis of $\text{im } P$ and a basis of $\ker P$ and combine them. In this basis the matrix of P has an identity block on the image and a zero block on the kernel. Its trace is therefore the dimension of the image, which is $\text{rank}(P)$. Since trace and rank are invariant under change of basis, $\text{tr}(P) = \text{rank}(P)$ in every basis.

Putnam 2006 B4

The maximum is 2^k . It is achieved by taking V to be the span of the first k coordinate vectors; then $V \cap Z$ consists of all 0-1 vectors supported in the first k coordinates.

To prove that no larger value is possible, choose k coordinate positions such that projection onto those coordinates is injective on V . Such coordinates exist by row-reducing a basis matrix for V and taking pivot columns. If two vectors in V agree on those pivot coordinates, their difference is a vector in V whose pivot coordinates are all zero, so the row-reduced description forces the difference to be zero. Therefore each vector in $V \cap Z$ is determined by its k pivot coordinates, and there are at most 2^k possible 0-1 patterns on those coordinates. Hence $|V \cap Z| \leq 2^k$.

Section 16: Sequences and Series

Section 16, Exercise 3

The recurrence $a_{n+1} = (a_n + 2/a_n)/2$ is Newton's method for $x^2 = 2$. If $a_n > \sqrt{2}$, then

$$a_{n+1} - \sqrt{2} = \frac{a_n + 2/a_n - 2\sqrt{2}}{2} = \frac{(a_n - \sqrt{2})^2}{2a_n} > 0,$$

so the next term is still above $\sqrt{2}$. Also $a_{n+1} < a_n$ whenever $a_n > \sqrt{2}$, because this inequality is equivalent to $2/a_n < a_n$. If $a_1 < \sqrt{2}$, then $a_2 > \sqrt{2}$ by the same square formula. Hence from the second term onward the sequence is decreasing and bounded below by $\sqrt{2}$, so it converges. Passing to the limit L gives $L = (L + 2/L)/2$, hence $L = \sqrt{2}$.

Putnam 2006 B6

The sequence is increasing. It cannot be bounded, because if $a_n \leq M$ forever, then the increments $a_n^{-1/k}$ would be at least $M^{-1/k}$ and the sequence would eventually exceed M . Hence $a_n \rightarrow \infty$.

Set $b_n = a_n^{(k+1)/k}$. By the mean value theorem,

$$b_{n+1} - b_n = \frac{k+1}{k} \xi_n^{1/k} (a_{n+1} - a_n)$$

for some ξ_n between a_n and a_{n+1} . Since $a_{n+1} - a_n = a_n^{-1/k}$, this becomes

$$b_{n+1} - b_n = \frac{k+1}{k} \left(\frac{\xi_n}{a_n} \right)^{1/k}.$$

Now $a_{n+1}/a_n = 1 + a_n^{-(k+1)/k} \rightarrow 1$, so $\xi_n/a_n \rightarrow 1$. Therefore $b_{n+1} - b_n \rightarrow (k+1)/k$. By the Cesaro mean theorem, $b_n/n \rightarrow (k+1)/k$. Since $b_n^k = a_n^{k+1}$, we get

$$\lim_{n \rightarrow \infty} \frac{a_n^{k+1}}{n^k} = \left(\frac{k+1}{k} \right)^k.$$

Section 17: Real Analysis

Section 17, Exercise 2

Define $g_i = f(i/n) - f((i-1)/n)$ for $i = 1, \dots, n$. If all g_i were positive, their sum would be positive; if all were negative, their sum would be negative. But

$$\sum_{i=1}^n g_i = f(1) - f(0) = 0.$$

Thus either some $g_i = 0$, giving the desired equality with $x = (i-1)/n$, or there are both positive and negative increments. For the continuous version on all subintervals, define $h(x) = f(x+1/n) - f(x)$ on $[0, 1-1/n]$. If h were always positive, summing $h(0), h(1/n), \dots, h((n-1)/n)$ would contradict $f(1) = f(0)$; the same holds if h were always negative. By continuity, h has a zero.

Putnam 2004 A1

Let $D(N) = 5S(N) - 4N$. The condition $S(N) < 0.8N$ is equivalent to $D(N) < 0$, and $S(N) > 0.8N$ is equivalent to $D(N) > 0$. After one more shot, D changes by 1 if the shot is made, because $5(S+1) - 4(N+1) = D+1$, and changes by -4 if the shot is missed.

Suppose D is negative at some time and positive later. At the first time it is nonnegative, the preceding value was negative. The last change cannot have been a miss, since a miss decreases D . It must have been a made shot, which increases D by exactly 1. The only way to move from a negative integer to a nonnegative integer by adding 1 is to move from -1 to 0 . Therefore at that time $D(N) = 0$, meaning $S(N) = 0.8N$. The statement is a discrete intermediate value theorem with unequal jump sizes.

Section 18: Calculus and Functional Methods

Section 18, Exercise 5

Since $f(x+y) = f(x)f(y)$ and f is positive, define $g(x) = \log f(x)$. Then $g(x+y) = g(x) + g(y)$ and g is continuous. The standard continuous Cauchy equation proof gives $g(q) = qg(1)$ for rational q , first for integers by repeated addition and then for rational numbers by dividing. For real x , choose rationals $q_j \rightarrow x$. Continuity gives $g(x) = \lim g(q_j) = xg(1)$. Therefore $f(x) = e^{cx}$ with $c = g(1)$. Conversely every function of the form e^{cx} works.

Putnam 2002 B3

Put $t = 1/n$, so $0 < t \leq 1/2$. We need to estimate

$$\frac{1}{e} - (1-t)^{1/t} = \frac{1}{e} (1 - e^{-u}),$$

where

$$u = -1 - \frac{1}{t} \log(1-t) = \frac{t}{2} + \frac{t^2}{3} + \frac{t^3}{4} + \cdots.$$

For the upper bound, the series gives $u < t$ when $0 < t \leq 1/2$, and $1 - e^{-u} < u$, so $1 - e^{-u} < t$. Thus the difference is less than $t/e = 1/(ne)$.

For the lower bound, it is enough to prove $e^{-u} < 1 - t/2$. Taking logarithms, this is equivalent to

$$1 + \frac{1}{t} \log(1-t) < \log(1-t/2).$$

Using the logarithm series, the left side is $-t/2 - t^2/3 - t^3/4 - \cdots$, while the right side is $-t/2 - t^2/8 - t^3/24 - \cdots$. Starting with the t^2 term, each coefficient on the left is more negative, so the inequality holds. Hence $1 - e^{-u} > t/2$, and the desired lower bound is $1/(2ne)$.

Section 19: Polynomials and Complex Numbers**Section 19, Exercise 6**

Let $\omega = e^{2\pi i/3}$. If $P(x) = (1+x)^n = \sum c_j x^j$, then the sum of the coefficients with exponents congruent to 1 modulo 3 is

$$\frac{P(1) + \omega^{-1}P(\omega) + \omega^{-2}P(\omega^2)}{3}.$$

This is the roots-of-unity filter: the coefficient c_j is multiplied by $1 + \omega^{j-1} + \omega^{2(j-1)}$, which equals 3 when $j \equiv 1 \pmod{3}$ and equals 0 otherwise. Substituting $P(x) = (1+x)^n$ gives the desired closed form. The method works whenever exponents are selected by a congruence class.

Putnam 2005 A3

If a zero of g' is also a zero of p , then it already has absolute value 1. Otherwise let the zeros of p be $\zeta_1, \dots, \zeta_n$, each with $|\zeta_j| = 1$. At a point z that is not a zero of p , the logarithmic derivative gives

$$\frac{g'(z)}{g(z)} = \sum_{j=1}^n \frac{1}{z - \zeta_j} - \frac{n}{2z}.$$

If $g'(z) = 0$, then after multiplying by z we get

$$\sum_{j=1}^n \frac{z}{z - \zeta_j} = \frac{n}{2}.$$

For $|\zeta| = 1$, a direct calculation shows

$$\operatorname{Re} \left(\frac{z}{z - \zeta} \right) = \frac{|z|^2 - \operatorname{Re}(z\bar{\zeta})}{|z - \zeta|^2}.$$

Subtracting $1/2$ from this real part gives

$$\frac{|z|^2 - 1}{2|z - \zeta|^2}.$$

Thus every term has real part greater than $1/2$ if $|z| > 1$, and less than $1/2$ if $|z| < 1$. The sum can have real part exactly $n/2$ only when $|z| = 1$. Hence every zero of g' has absolute value 1.

Section 20: Functional Equations and Synthesis

Section 20, Exercise 4

Because S is nonempty, choose $a \in S$. Closure gives $0 = |a - a| \in S$. If S has no positive element, then $S = \{0\}$.

Otherwise let m be the least positive element of S . We first prove that every element of S is divisible by m . Given $a \in S$, write $a = tm + r$ with $0 \leq r < m$. Repeatedly taking the absolute difference with m shows that

$$a - m, a - 2m, \dots, a - tm = r$$

all lie in S as long as the displayed numbers are nonnegative. If $r > 0$, this contradicts the minimality of m . Hence $r = 0$.

Let M be the largest element of S . We have $M = qm$ for some positive integer q . Starting with M and repeatedly taking the difference with m shows that

$$(q - 1)m, (q - 2)m, \dots, m, 0$$

all belong to S . Conversely, the divisibility argument and maximality of M show that no other elements can occur. Therefore $S = \{0, m, 2m, \dots, qm\}$. The proof is the Euclidean algorithm expressed as a closure argument.

Putnam 2001 A1

The given identity says $(x * y) * x = y$ for every pair x, y . Apply this identity first with $x = b$ and $y = a$. It gives

$$(b * a) * b = a.$$

Now apply the same identity again, this time with $x = b * a$ and $y = b$. Then

$$((b * a) * b) * (b * a) = b.$$

Using the previous displayed equality, the left side becomes $a * (b * a)$. Hence $a * (b * a) = b$, as required. The proof is a good example of working backward in an algebraic structure: the target contains $b * a$, so we make $b * a$ the outer variable in the defining identity.

Appendix Exercises

Appendix A, Exercise 6

A useful theorem sheet is organized by triggers rather than by names. For Cauchy-Schwarz, write entries such as “sum of fractions with squared numerators,” “dot product bounded by norms,” and “need to turn two sums into one square.” For compactness, write “continuous function on closed bounded interval,” “finite-dimensional norm comparison,” and “choose an extremal object.” This format makes the sheet useful during problem solving because the problem usually shows a trigger before it shows a theorem name.

Appendix B, Exercise 1

The flaw is division by $a - b$. From $a^2 = b^2$ one may factor $(a - b)(a + b) = 0$, but if $a = b$, then $a - b = 0$, and division by zero is invalid. The correct conclusion is that either $a = b$ or $a = -b$. This is a common proof-writing error: an operation is performed before checking that its hypothesis, here nonzero divisor, is satisfied.

Appendix C, Exercise 5

For a ninety-minute simulation, reconstruct the session in intervals rather than writing only “I ran out of time.” Record when the first survey ended, when the main idea for each attempted problem appeared, when the clean write-up began, and what occupied the final ten minutes. Then classify each loss:

- a *recognition loss* means a known method was not identified;
- an *execution loss* means the right method led to unfinished algebra or an unproved lemma;
- a *triage loss* means too long was spent without a new fact;
- a *conversion loss* means a correct idea was found but clean writing began too late.

Finish with one measurable adjustment for the next packet, such as “complete the survey by minute six” or “begin the final write-up no later than minute sixty-five.” The postmortem is useful only when it changes the next rehearsal.

Appendix D, Exercise 3

When recording modular arithmetic problems, do not write only the modulus. Also write why that modulus was natural. A parity modulus may be forced by squares or by alternating signs. A modulus such as p may be forced by a prime in the statement. A modulus such as n may be forced by a desired divisibility conclusion. The useful archive entry is “modulus m because it makes the move invariant” or “modulus p^k because valuations matter,” not simply “used mod m .”

§F Additional Worked Exercise Solutions

The following pages add more complete solutions to the course exercises. They are included to make the manuscript useful not only as a list of techniques, but also as a model for how those techniques are written under contest conditions.

More solutions for Section 1

Section 1, Exercise 2

Let m and M be the minimum and maximum elements of the finite set S . Since S is closed under averages, $(m + M)/2$ lies in S . If $m < M$, then $m < (m + M)/2 < M$, so the new element is strictly between the minimum and maximum. That alone is not a contradiction, but repeating the idea with m and $(m + M)/2$ gives another element, then another, producing infinitely many distinct elements between m and M . A finite set cannot contain infinitely many distinct elements, so $m = M$. Therefore S has only one element.

Section 1, Exercise 5

Assume such a function exists. Starting from any n , the equation gives $f(f(n)) = n + 1$. Applying f to both sides in the same pattern yields $f(f(f(n))) = f(n + 1)$. But also replacing n by $f(n)$ gives $f(f(f(n))) = f(n) + 1$. Hence $f(n + 1) = f(n) + 1$ for all n . By induction, $f(n) = f(1) + n - 1$. Substituting into $f(f(n)) = n + 1$ gives

$$f(1) + f(n) - 1 = n + 1,$$

so $f(1) + f(1) + n - 2 = n + 1$. Thus $2f(1) = 3$, impossible in integers. The contradiction comes from turning the second iterate condition into a first-order recurrence.

More solutions for Section 2

Section 2, Exercise 2

We prove by strong induction that every integer $n \geq 2$ is a product of primes. If n is prime, it is already a product of one prime. If n is composite, then $n = ab$ with $2 \leq a < n$ and $2 \leq b < n$. By the induction hypothesis, both a and b are products of primes. Multiplying those prime factorizations gives a prime factorization of n . Strong induction is natural here because the factors a and b are smaller than n , but their sizes are not predetermined.

Section 2, Exercise 3

Suppose there are only finitely many primes, say $p_1, \dots, p_k$. Consider $N = p_1 p_2 \cdots p_k + 1$. The integer N is greater than 1, so it has a prime divisor q . This q must be one of the listed primes. But every listed prime divides the product $p_1 \cdots p_k$, so it leaves remainder 1 when dividing N . Thus none of the listed primes divides N , contradiction. The proof works because the constructed number is designed to avoid every prime in the supposed complete list.

More solutions for Section 3

Section 3, Exercise 3

The polynomial $P(x) - x^2$ has degree at most 3 and vanishes at $x = 0, 1, 2, 3$. A nonzero polynomial of degree at most 3 cannot have four distinct roots. Therefore $P(x) - x^2$ is identically zero, and $P(x) = x^2$ for all x . Hence $P(4) = 16$. The useful move is to compare the unknown polynomial with the obvious polynomial that matches the data.

Section 3, Exercise 4

The inequality $x^4 + y^4 \geq 2x^2y^2$ is equivalent to

$$x^4 - 2x^2y^2 + y^4 \geq 0,$$

which factors as $(x^2 - y^2)^2 \geq 0$. Equality holds exactly when $x^2 = y^2$, that is, when $x = y$ or $x = -y$. The important detail is that the equality condition is broader than $x = y$, because the variables appear only through their squares.

More solutions for Section 4

Section 4, Exercise 3

By Cauchy-Schwarz in Engel form,

$$\frac{1}{a} + \frac{1}{b} + \frac{1}{c} = \frac{1^2}{a} + \frac{1^2}{b} + \frac{1^2}{c} \geq \frac{(1+1+1)^2}{a+b+c} = \frac{9}{a+b+c}.$$

Equality occurs when $a = b = c$. This is one of the cleanest examples of Engel form: a sum of reciprocal terms becomes a square of a sum divided by the sum of denominators.

Section 4, Exercise 6

Let the fixed sum be S , with $0 < S < n\pi$. Since $\log(\sin x)$ is concave on $(0, \pi)$, Jensen's inequality gives

$$\frac{1}{n} \sum_{i=1}^n \log(\sin x_i) \leq \log\left(\sin \frac{S}{n}\right).$$

Exponentiating yields

$$\prod_{i=1}^n \sin x_i \leq \left(\sin \frac{S}{n}\right)^n.$$

The maximum is attained when $x_1 = \cdots = x_n = S/n$. The concavity condition is essential, and it holds on the entire interval $(0, \pi)$.

More solutions for Section 5

Section 5, Exercise 1

The inequality $\ln x \geq 1 - 1/x$ can be proved by applying the already known inequality $\ln y \leq y - 1$ to $y = 1/x$. This gives $\ln(1/x) \leq 1/x - 1$, or $-\ln x \leq 1/x - 1$. Multiplying by -1 gives $\ln x \geq 1 - 1/x$.

Equality occurs at $x = 1$. This solution shows how a tangent-line inequality can be reflected by replacing x with its reciprocal.

Section 5, Exercise 4

For $x > 0$, direct algebra gives

$$x + \frac{1}{x} - 2 = \frac{x^2 + 1 - 2x}{x} = \frac{(x-1)^2}{x}.$$

The right side is nonnegative, so $x + 1/x \geq 2$, with equality exactly when $x = 1$. If $1/2 \leq x \leq 2$, then $1/x \geq 1/2$. Therefore

$$\frac{(x-1)^2}{x} \geq \frac{1}{2}(x-1)^2,$$

and substitution in the identity yields

$$x + \frac{1}{x} \geq 2 + \frac{1}{2}(x-1)^2.$$

The exact identity is stronger than a Taylor approximation: it displays both the equality case and the size of the error for every positive x .

More solutions for Section 6

Section 6, Exercise 1

Every integer is congruent modulo 8 to one of 0, 1, 2, 3, 4, 5, 6, 7. Squaring these residues gives

$$0, 1, 4, 1, 0, 1, 4, 1 \pmod{8}.$$

Thus the quadratic residues modulo 8 are 0, 1, 4. A shorter proof observes that even squares are 0 or 4 modulo 8, while odd squares are 1 modulo 8.

Section 6, Exercise 5

We solve $x \equiv 2 \pmod{5}$ and $x \equiv 3 \pmod{7}$. Write $x = 2 + 5t$. Then $2 + 5t \equiv 3 \pmod{7}$, so $5t \equiv 1 \pmod{7}$. Since $5^{-1} \equiv 3 \pmod{7}$, we get $t \equiv 3 \pmod{7}$. Thus $x = 2 + 5(3 + 7k) = 17 + 35k$. The solution is unique modulo 35 because two solutions would have difference divisible by both 5 and 7, hence by 35.

More solutions for Section 7

Section 7, Exercise 1

The equation $x^2 - y^2 = 15$ factors as $(x-y)(x+y) = 15$. Both factors have the same parity, and since their product is odd, both are odd. The integer factor pairs of 15 are (1, 15), (3, 5), and their sign variants. Solving $x - y = d$ and $x + y = e$ gives $x = (d + e)/2$ and $y = (e - d)/2$. The positive factor pairs yield $(x, y) = (8, 7)$ and $(4, 1)$, and sign choices give the corresponding negative solutions. The factorization turns the quadratic equation into a finite divisor problem.

Section 7, Exercise 4

If $x^2 = 2y^2$, then x^2 is even, so x is even. Write $x = 2z$. Then $4z^2 = 2y^2$, so $y^2 = 2z^2$, and y is even. Thus any nonzero integer solution produces a smaller nonzero solution by dividing both x and y by 2. Repeating gives an infinite descent, impossible in the positive integers. Therefore the only integer solution is $x = y = 0$.

More solutions for Section 8

Section 8, Exercise 3

The integers from 1 to p^a that are not relatively prime to p^a are exactly the multiples of p . There are p^{a-1} such multiples. Hence

$$\varphi(p^a) = p^a - p^{a-1}.$$

The formula counts units modulo a prime power by subtracting the only possible obstruction, divisibility by the prime p .

Section 8, Exercise 5

If $\varphi(n) = n/2$, then exactly half of the residues from 1 to n are coprime to n . For $n = 2^a$, this is true because $\varphi(2^a) = 2^a - 2^{a-1} = 2^{a-1}$. If n has an odd prime divisor p , then the product formula

$$\frac{\varphi(n)}{n} = \prod_{q|n} \left(1 - \frac{1}{q}\right)$$

contains a factor $(1 - 1/p) < 1$ in addition to the possible factor $1/2$ from $q = 2$, so the product is less than $1/2$ if $2 \mid n$, and is not equal to $1/2$ if $2 \nmid n$ because all factors have odd-prime denominators. Therefore the solutions are exactly $n = 2^a$ with $a \geq 1$.

More solutions for Section 9

Section 9, Exercise 2

Count the number of ways to choose n people from two disjoint teams of n people each. Directly, the answer is $\binom{2n}{n}$. If exactly k people are chosen from the first team, then $n - k$ are chosen from the second team, giving $\binom{n}{k} \binom{n}{n-k} = \binom{n}{k}^2$ choices. Summing over k gives

$$\sum_{k=0}^n \binom{n}{k}^2 = \binom{2n}{n}.$$

The identity is not an algebraic coincidence; it records two descriptions of the same selection.

Section 9, Exercise 5

Count incidences (x, A_i) where x is an element of the ground set and $x \in A_i$. Counting by subsets gives mr , because there are m subsets and each has r elements. Counting by elements gives ns , because there are n elements and each lies in exactly s subsets. Since both counts describe the same incidence set, $mr = ns$.

More solutions for Section 10

Section 10, Exercise 2

Reduce the $n + 1$ integers modulo n . There are only n residue classes. By the pigeonhole principle, two of the integers have the same residue modulo n . Their difference is divisible by n . This is the most basic residue-class pigeonhole argument, and it reappears in partial-sum form throughout the Putnam.

Section 10, Exercise 4

In a tournament on n vertices, every pair of vertices contributes exactly one directed edge. Therefore the total outdegree is the total number of directed edges, which is $\binom{n}{2}$. The average outdegree is

$$\frac{\binom{n}{2}}{n} = \frac{n-1}{2}.$$

At least one vertex has outdegree at least the average. This proof is an averaging argument disguised as a graph statement.

More solutions for Section 11

Section 11, Exercise 2

The generating function for $a_n = n$ with $n \geq 0$ is

$$\sum_{n \geq 0} nx^n.$$

Start from the geometric series $\sum_{n \geq 0} x^n = 1/(1-x)$ and differentiate to get $\sum_{n \geq 1} nx^{n-1} = 1/(1-x)^2$. Multiplying by x gives

$$\sum_{n \geq 0} nx^n = \frac{x}{(1-x)^2}.$$

The differentiation is valid as a formal power series, so convergence questions are not needed for the algebraic identity.

Section 11, Exercise 6

Put $s_n = (x_n, x_{n+1}) \in R^2$. The recurrence is equivalent to the first-order state equation

$$s_{n+1} = T(s_n), \quad T(u, v) = (v, F(u, v)).$$

If $|R| = q$, then there are only q^2 possible states. Among $s_0, s_1, \dots, s_{q^2}$ two must agree, say $s_i = s_j$ with $0 \leq i < j \leq q^2$. Applying T repeatedly gives $s_{i+r} = s_{j+r}$ for every $r \geq 0$. Hence the state sequence, and therefore the original sequence, is periodic from index i onward.

Now suppose T is bijective. It has an inverse on the finite set R^2 . Apply T^{-i} to $s_i = s_j$ to obtain

$$s_0 = s_{j-i}.$$

Applying further powers of T gives $s_n = s_{n+j-i}$ for every $n \geq 0$. Thus the sequence is periodic from its initial state. Bijectivity eliminates the transient tail because every state has a unique predecessor.

More solutions for Section 12

Section 12, Exercise 2

Let T be a finite tree with at least two vertices. Take a longest simple path, say from u to v . If u had degree at least 2, it would have a neighbor other than the next vertex on the path. That neighbor cannot already occur later on the path, because the resulting extra edge would create a cycle. Therefore it could be placed before u to extend the path, contradicting maximality. Hence u has degree 1. The same argument applies to v , and $u \neq v$ because the path contains at least one edge. Thus T has at least two distinct leaves.

Section 12, Exercise 5

Repeatedly delete any vertex of degree at most k . If the process deletes every vertex, then at the moment each vertex was deleted it had degree at most k , so the total number of edges removed is at most kn . This would imply average degree at most $2k$, contrary to the hypothesis. Hence some nonempty subgraph remains. In that remaining subgraph no vertex has degree at most k , so the minimum degree is at least $k + 1$.

More solutions for Section 13

Section 13, Exercise 1

The perpendicular bisector of AB is the set of points equidistant from A and B . The perpendicular bisector of AC is the set of points equidistant from A and C . Their intersection point O is therefore equidistant from A , B , and C . Hence O also lies on the perpendicular bisector of BC . Thus all three perpendicular bisectors meet at one point, the circumcenter.

Section 13, Exercise 3

Let $ABCD$ be cyclic. The inscribed angle theorem says that $\angle A$ intercepts arc BCD , while $\angle C$ intercepts the complementary arc DAB . The measures of these two arcs add to 360° , so the corresponding inscribed angles add to 180° . Therefore opposite angles in a cyclic quadrilateral are supplementary.

More solutions for Section 14

Section 14, Exercise 1

Expand both squared norms using dot products:

$$\|u + v\|^2 = (u + v) \cdot (u + v) = \|u\|^2 + 2u \cdot v + \|v\|^2,$$

$$\|u - v\|^2 = (u - v) \cdot (u - v) = \|u\|^2 - 2u \cdot v + \|v\|^2.$$

Adding cancels the cross terms and gives the parallelogram identity. The cancellation reflects the geometry that the two diagonals of a parallelogram balance each other.

Section 14, Exercise 5

Let the fixed point be A . For P on the unit circle centered at the origin, the distance $|P - A|$ is maximized when P points directly away from A . Algebraically,

$$|P - A|^2 = |P|^2 + |A|^2 - 2P \cdot A = 1 + |A|^2 - 2P \cdot A.$$

This is largest when $P \cdot A$ is as small as possible. If $A \neq 0$, the minimum is $-|A|$, attained at $P = -A/|A|$, giving maximum distance $1 + |A|$. If $A = 0$, every point on the unit circle gives distance 1.

More solutions for Section 15

Section 15, Exercise 2

The all-ones matrix J_n sends a vector v to a vector all of whose entries equal the sum of the entries of v . The vector $(1, 1, \dots, 1)$ is therefore an eigenvector with eigenvalue n . Any vector whose coordinates sum to 0 is sent to the zero vector, so it is an eigenvector with eigenvalue 0. The subspace of coordinate-sum-zero vectors has dimension $n - 1$, so the eigenvalues are n once and 0 with multiplicity $n - 1$.

Section 15, Exercise 4

If $AB = I$ for square matrices, then the linear map B is injective: if $Bx = 0$, then $x = ABx = 0$. An injective linear map from a finite-dimensional vector space to itself is also surjective. Hence for every y there is an x with $Bx = y$. Then $BAy = BABx = BIx = Bx = y$. Since BA acts as the identity on every vector, $BA = I$.

More solutions for Section 16

Section 16, Exercise 1

For $n \geq 1$, compare the decreasing function $1/x^2$ with the series. The integral test gives

$$\sum_{n=2}^{\infty} \frac{1}{n^2} \leq \int_1^{\infty} \frac{dx}{x^2} = 1.$$

Adding the first term shows convergence. The comparison is effective because $1/x^2$ has a finite tail integral, unlike $1/x$.

Section 16, Exercise 4

Group the harmonic sum dyadically:

$$1 + \frac{1}{2} + \left(\frac{1}{3} + \frac{1}{4}\right) + \left(\frac{1}{5} + \dots + \frac{1}{8}\right) + \dots.$$

For the upper bound up to N , compare with the integral or note that the block from $2^j + 1$ to 2^{j+1} has at most 2^j terms, each at most $1/2^j$, so each block contributes at most 1. There are $O(\log N)$ blocks. Hence $\sum_{n=1}^N 1/n = O(\log N)$.

More solutions for Section 17

Section 17, Exercise 1

Let $p(x)$ be a real polynomial of odd degree with leading coefficient c . As $x \rightarrow \infty$, the leading term dominates, so $p(x)$ has the sign of c . As $x \rightarrow -\infty$, the leading term has the opposite sign because the degree is odd. Thus p takes both positive and negative values for sufficiently large positive and negative inputs. By continuity, the intermediate value theorem gives a real root.

Section 17, Exercise 5

Suppose $\lim_{x \rightarrow \infty} f(x) = L_+$ and $\lim_{x \rightarrow -\infty} f(x) = L_-$ are finite. Choose R so large that $|f(x)| \leq |L_+| + 1$ for $x \geq R$ and $|f(x)| \leq |L_-| + 1$ for $x \leq -R$. On the compact interval $[-R, R]$, the continuous function f is bounded by the extreme value theorem. Combining the two tails with this compact interval proves that f is bounded on all of $\mathbb{R}$.

More solutions for Section 18

Section 18, Exercise 1

Let $h(x) = e^x - 1 - x$. Then $h'(x) = e^x - 1$ and $h''(x) = e^x > 0$, so h' is increasing and vanishes at $x = 0$. Hence h decreases on $(-\infty, 0]$ and increases on $[0, \infty)$, with minimum $h(0) = 0$. Therefore $e^x \geq 1 + x$ for all real x . The derivative proof also identifies the equality case.

Section 18, Exercise 6

Set $y = 0$ in $f(x + y) = f(x) + f(y) + xy$. This gives $f(x) = f(x) + f(0)$, so $f(0) = 0$. Define $g(x) = f(x) - x^2/2$. Then

$$g(x + y) = f(x + y) - \frac{(x + y)^2}{2} = g(x) + g(y).$$

Since f is differentiable, g is differentiable, and every differentiable additive function is linear. Thus $g(x) = cx$, and $f(x) = x^2/2 + cx$. Conversely, every function of this form satisfies the equation.

More solutions for Section 19

Section 19, Exercise 1

If P has real coefficients and $P(z) = 0$, then conjugating the equation gives $\overline{P(z)} = 0$. Since the coefficients are real, $\overline{P(z)} = P(\bar{z})$. Therefore $P(\bar{z}) = 0$. Complex roots of real polynomials occur in conjugate pairs, with multiplicity as well if the factorization is tracked carefully.

Section 19, Exercise 4

The identity

$$x^n - y^n = (x - y)(x^{n-1} + x^{n-2}y + \cdots + xy^{n-2} + y^{n-1})$$

is verified by multiplying out: all middle terms cancel telescopically, leaving $x^n - y^n$. Thus $x - y$ divides $x^n - y^n$ in the polynomial ring, and in particular for every pair of integers x, y .

More solutions for Section 20

Section 20, Exercise 1

Setting $y = 0$ gives $f(x) = f(x) + f(0)$, so $f(0) = 0$. For positive integers, repeated addition gives $f(n) = nf(1)$. Also $0 = f(n + (-n)) = f(n) + f(-n)$, so $f(-n) = -f(n)$. Therefore every solution has the form $f(n) = cn$ for an integer constant $c = f(1)$. Conversely, every function $f(n) = cn$ satisfies the equation.

Section 20, Exercise 3

Represent red by 1 and blue by 0 in the vector space $\mathbb{F}_2^V$, where V is the vertex set. A move at vertex v adds the vector that has 1 at v and at every neighbor of v , and 0 elsewhere. Reachability from one coloring to another is therefore the question of whether their difference lies in the span of these move vectors. Equivalently, if A is the adjacency matrix and I the identity matrix over $\mathbb{F}_2$, the move vectors are the columns of $A + I$, so the reachable differences are the column space of $A + I$.

Appendix A, Exercise 3

Bezout's identity says that if $\gcd(a, b) = 1$, then there exist integers u, v such that $ua + vb = 1$. If a prime p divides ab and $p \nmid a$, then $\gcd(p, a) = 1$, so there exist integers u, v with $up + va = 1$. Multiplying by b gives $upb + vab = b$. Both terms on the left are divisible by p , so p divides b . This proves Euclid's lemma.

Appendix B, Exercise 2

The sequence $a_n = (-1)^n$ is bounded because $|a_n| = 1$ for every n , but it does not converge. The subsequence with even n is constantly 1, while the subsequence with odd n is constantly -1 . A convergent sequence has all subsequences converging to the same limit, so this bounded sequence is not convergent.

Appendix C, Exercise 2

A messy solution usually contains the right ideas in the wrong order. Start the rewrite by identifying the first sentence that makes the solution possible. Move definitions before they are used, move computations after the lemma that motivates them, and delete any calculation that does not feed the final conclusion. The finished proof should read as if the key idea was natural, even if it was discovered only after experimentation.

Appendix D, Exercise 5

When re-solving an old problem after two weeks, avoid looking at the original solution until you have written a complete attempt. If the proof returns quickly, the problem has become part of your working memory. If only the first idea returns, write down exactly where the reconstruction fails. That failure point identifies the lemma or transformation that still needs review.

§G Further Proof Models

This final supplement is meant to make proof style more explicit. It is organized by the mathematical section to which each model belongs, but a proof model is not presented as the solution to a numbered exercise unless it actually solves that exercise. Many Putnam solutions become short only after the solver has learned which reductions are safe. In the paragraphs below, the emphasis is on explaining why the chosen reduction is natural, which quantities are being preserved, and where the decisive estimate or structural observation enters.

Expanded solutions to revised exercises

Section 1, Exercise 1

Let $r \neq 1$ be reachable. Because r is not the initial value, it was produced by a final move from some earlier reachable integer a . There are exactly two cases.

If the final move was $a \mapsto 2a + 1$, then $r = 2a + 1$, so r is odd. If the final move was $a \mapsto 3a + 1$, then

$$r = 3a + 1 \equiv 1 \pmod{3}.$$

Thus every reachable value other than 1 lies in the union of the odd integers and the residue class 1 modulo 3. Notice the direction of the assertion: the last-move argument gives a necessary condition. It does not claim that every integer satisfying one of these congruence conditions is reachable.

Section 1, Exercise 3

Work modulo 2 and let B denote the number of black squares. Flipping one square changes B by 1 modulo 2, regardless of whether the square changes from black to white or from white to black.

A move first flips the n squares of a row and then the m squares of a column. The intersection square is flipped twice, so its net color is unchanged. Equivalently, exactly

$$(n - 1) + (m - 1) = m + n - 2$$

squares have their final colors reversed. Hence

$$B_{\text{after}} - B_{\text{before}} \equiv m + n - 2 \equiv m + n \pmod{2}.$$

When $m + n$ is even, this difference is 0 modulo 2. Therefore the parity of the number of black squares is invariant under every legal move.

Section 1, Exercise 6

Since

$$\frac{x_1 + \cdots + x_n}{n} = 0,$$

not every x_i can be negative. Choose an index j for which $x_j = \max_i x_i$. Then $x_j \geq 0$. The singleton set $\{x_j\}$ is nonempty, and because $n \geq 2$ it is a proper subset of the n indexed terms. Its average is $x_j \geq 0$, which proves the claim. The hypothesis $n \geq 2$ is necessary: for $n = 1$ there is no nonempty proper subset.

Section 3, Exercise 6

List the roots with multiplicity as $r_1, \dots, r_d$. Since P is monic,

$$P(x) = \prod_{i=1}^d (x - r_i).$$

At the integer k this becomes

$$P(k) = \prod_{i=1}^d (k - r_i).$$

Every factor is an integer. A product of integers has absolute value 1 only when every factor has absolute value 1. Thus $k - r_i \in \{-1, 1\}$ for every i , or equivalently

$$r_i \in \{k - 1, k + 1\}.$$

Conversely, suppose $k - 1$ has multiplicity u and $k + 1$ has multiplicity v , where $u + v = d$. Then

$$P(k) = (k - (k - 1))^u (k - (k + 1))^v = 1^u (-1)^v = (-1)^v.$$

Therefore $P(k) = 1$ exactly when the multiplicity of $k + 1$ is even, and $P(k) = -1$ exactly when that multiplicity is odd.

Section 5, Exercise 3

Expanding and using $abc = 1$ gives

$$\begin{aligned} (a - 1)(b - 1)(c - 1) &= abc - ab - ac - bc + a + b + c - 1 \\ &= a + b + c - ab - bc - ca. \end{aligned}$$

Consequently the proposed inequality holds exactly when $(a - 1)(b - 1)(c - 1) \geq 0$.

If one variable equals 1, the product on the right is 0, so equality holds. Suppose no variable equals 1. Since $abc = 1$, the three variables cannot all exceed 1 and cannot all be less than 1. Thus either exactly one is greater than 1, in which case two of the factors $a - 1, b - 1, c - 1$ are negative and their product is positive, or exactly two are greater than 1, in which case the product is negative. It follows that the inequality holds precisely when $a = b = c = 1$ or exactly one of a, b, c is greater than 1.

Section 5, Exercise 5

Put

$$A = x^2 + x + 1, \quad B = y^2 + y + 1, \quad C = z^2 + z + 1.$$

All three denominators are positive. Bringing the left side minus 1 over the common denominator ABC and simplifying with $xyz = 1$ gives

$$ABC \left(\frac{x^2}{A} + \frac{y^2}{B} + \frac{z^2}{C} - 1 \right) = x^2 y^2 + y^2 z^2 + z^2 x^2 - (x + y + z).$$

The last expression has a useful sum-of-squares form. Since $xyz = 1$,

$$\begin{aligned} x^2 y^2 + y^2 z^2 + z^2 x^2 - (x + y + z) \\ = \frac{1}{2} ((xy - yz)^2 + (yz - zx)^2 + (zx - xy)^2) \geq 0. \end{aligned}$$

Indeed, expanding the three squares produces the three square terms and subtracts

$$xy \cdot yz + yz \cdot zx + zx \cdot xy = xyz(x + y + z) = x + y + z.$$

Division by the positive number ABC proves the inequality. Equality requires $xy = yz = zx$; positivity gives $x = y = z$, and $xyz = 1$ then gives $x = y = z = 1$.

Section 7, Exercise 6

The quotient k is a positive integer and the defining equation is

$$(1) \quad x^2 + y^2 = k(xy + 1).$$

Assume, for contradiction, that some solution has nonsquare k . Among all positive solutions for this fixed k , choose one with $x + y$ minimal, and relabel the variables so that $x \geq y$.

Regard (1) as a quadratic equation in x :

$$X^2 - kyX + (y^2 - k) = 0.$$

One integer root is x , so the other root

$$x' = ky - x$$

is also an integer. By the product of the roots,

$$(2) \quad xx' = y^2 - k.$$

We first prove $y^2 > k$. Equality would make $k = y^2$ a square. If $y^2 < k$, then (2) gives $x' < 0$, so $x - ky = -x' > 0$. Rearranging (2) yields

$$x(x - ky) = k - y^2.$$

The left side is at least x , while $x > ky \geq k$; the right side is strictly less than k . This is impossible. Hence $y^2 > k$.

Equation (2) now shows $x' > 0$. Moreover,

$$x' = \frac{y^2 - k}{x} < \frac{y^2}{x} \leq y,$$

because $x \geq y$. Vieta's formulas show that (x', y) satisfies the same equation (1), but

$$x' + y < x + y.$$

This contradicts the minimal choice of (x, y) . Therefore a nonsquare quotient cannot occur, and k must be a perfect square.

Section 8, Exercise 6

If $n = p$ is prime, its positive divisors are only 1 and p . The equation $\sum_{d|p} d = 2p$ would then give

$$1 + p = 2p,$$

so $p = 1$, which is not prime. Hence a number satisfying the equation cannot be prime.

For the requested examples,

$$1 + 2 + 3 + 6 = 12 = 2 \cdot 6$$

and

$$1 + 2 + 4 + 7 + 14 + 28 = 56 = 2 \cdot 28.$$

Finally, $496 = 2^4 \cdot 31$. Since the two prime powers are coprime, the divisor sum factors:

$$\sum_{d|496} d = (1 + 2 + 2^2 + 2^3 + 2^4)(1 + 31) = 31 \cdot 32 = 992 = 2 \cdot 496.$$

Thus 6, 28, and 496 all satisfy the equation.

Section 9, Exercise 6

First choose the line ℓ . Any unordered pair of points of S determines one line, so there are $\binom{n}{2}$ choices. Because no three points are collinear, exactly two points of S lie on ℓ . There are therefore 2 choices for the distinguished point $P \in S \cap \ell$ and $n - 2$ choices for $Q \in S \setminus \ell$. The multiplication principle gives

$$\binom{n}{2} \cdot 2 \cdot (n - 2) = n(n - 1)(n - 2)$$

ordered triples.

As a check, choose an ordered triple of distinct points (P, R, Q) . The first two determine $\ell = PR$, and the no-three-collinear hypothesis ensures $Q \notin \ell$. This gives a bijection with the triples counted above and produces the same number $n(n - 1)(n - 2)$.

Section 10, Exercise 6

Divide the unit square into an $m \times m$ grid of m^2 smaller squares, each of side length $1/m$. Assign boundary points consistently, for example by making cells half-open except along the outer boundary. This convention gives each point exactly one cell.

There are $m^2 + 1$ points and only m^2 cells, so the pigeonhole principle places two points in the same small square. The greatest possible distance between two points of that square is its diagonal:

$$\sqrt{\left(\frac{1}{m}\right)^2 + \left(\frac{1}{m}\right)^2} = \frac{\sqrt{2}}{m}.$$

The two selected points are therefore at distance at most $\sqrt{2}/m$.

Section 13, Exercise 4

Rotate the plane by 60° clockwise about A . Let B' and P' be the images of B and P . Rotation preserves distance, so

$$B'P' = BP.$$

Also $AP = AP'$ and $\angle PAP' = 60^\circ$, hence triangle APP' is equilateral and $P'P = AP$. Therefore the broken path $B' \rightarrow P' \rightarrow P \rightarrow C$ has length

$$B'P' + P'P + PC = PB + PA + PC.$$

By the triangle inequality for a broken path,

$$PA + PB + PC \geq B'C.$$

The rays AB' and AC make an angle of 120° , and $AB' = AC = s$. The law of cosines gives

$$B'C^2 = s^2 + s^2 - 2s^2 \cos 120^\circ = 3s^2.$$

Thus $B'C = \sqrt{3}s$, proving the desired inequality. At the center of the equilateral triangle, the points B', P', P, C lie on one segment in that order, so equality holds; equivalently each of PA, PB, PC equals $s/\sqrt{3}$.

Section 15, Exercise 6

Place the vectors $v_1, \dots, v_m$ as the columns of an $n \times m$ matrix A . Since $\text{rank}(A) \leq n < m$, the rank-nullity theorem gives

$$\dim \ker A = m - \text{rank}(A) > 0.$$

Choose a nonzero vector $(c_1, \dots, c_m)^T \in \ker A$. Then

$$c_1 v_1 + \dots + c_m v_m = 0,$$

which is the required nontrivial linear relation.

For the application, define

$$w_i = (1, t_i, t_i^2, \dots, t_i^n) \in \mathbb{R}^{n+1} \quad (1 \leq i \leq n+2).$$

There are $n+2$ vectors in an $(n+1)$ -dimensional space, so they satisfy a nontrivial relation

$$c_1 w_1 + \dots + c_{n+2} w_{n+2} = 0.$$

Comparing coordinate $j+1$ for $j = 0, \dots, n$ gives

$$\sum_{i=1}^{n+2} c_i t_i^j = 0.$$

The coefficients are not all zero because the vector relation was nontrivial.

Section 16, Exercise 6

The recurrence inequality gives

$$0 < a_{n+1} \leq a_n - a_n^2 = a_n(1 - a_n) < a_n.$$

Thus (a_n) is decreasing and bounded below by 0, so it converges to some $L \geq 0$. Taking limits in $a_{n+1} \leq a_n - a_n^2$ gives

$$L \leq L - L^2.$$

Hence $L^2 \leq 0$, so $L = 0$.

For the quantitative estimate, positivity permits reciprocals:

$$\frac{1}{a_{n+1}} \geq \frac{1}{a_n(1 - a_n)} = \frac{1}{a_n} + \frac{1}{1 - a_n} \geq \frac{1}{a_n} + 1.$$

Iterating from 1 to $n-1$ yields

$$\frac{1}{a_n} \geq \frac{1}{a_1} + n - 1.$$

Taking reciprocals, with both sides positive, gives

$$a_n \leq \frac{1}{n - 1 + 1/a_1}.$$

This bound proves convergence again and records its order of decay.

Section 17, Exercise 4

The identity $f(f(x)) = x$ implies injectivity: if $f(x) = f(y)$, applying f to both sides gives $x = y$. A continuous injective function on an interval is strictly monotone, so f is either strictly increasing or strictly decreasing.

Suppose f is strictly increasing. If $f(x) > x$ for some x , applying the increasing function f gives

$$f(f(x)) > f(x),$$

or $x > f(x)$, a contradiction. Similarly, if $f(x) < x$, then applying f gives $x < f(x)$, again a contradiction. Hence $f(x) = x$ for every x . If f is not increasing, the monotonicity dichotomy forces it to be strictly decreasing. These are exactly the two possibilities asserted.

Section 17, Exercise 6

For a fixed real number a , define

$$g_a(x) = f(x) - ax(1 - x).$$

The function g_a is continuous on $[0, 1]$ and differentiable on $(0, 1)$. Moreover,

$$g_a(0) = f(0) = 0, \quad g_a(1) = f(1) = 0.$$

Rolle's theorem therefore gives a point $c \in (0, 1)$ such that $g'_a(c) = 0$. Thus g_a has a critical point in the open interval for every real a . No additional condition on a is needed.

Section 20, Exercise 5

Fix n . Repeated use of the functional equation gives

$$f(n + qf(n)) = f(n) \quad (q = 0, 1, 2, \dots).$$

Indeed, the case $q = 0$ is immediate. If it holds for q , apply the given identity at the input $n + qf(n)$:

$$\begin{aligned} f(n + (q + 1)f(n)) &= f(n + qf(n) + f(n + qf(n))) \\ &= f(n + qf(n)) = f(n). \end{aligned}$$

Because $1 \leq f(n) \leq M$, the number $L = \text{lcm}(1, 2, \dots, M)$ is divisible by $f(n)$. Write $L = qf(n)$ for a positive integer q . The displayed identity gives

$$f(n + L) = f(n).$$

This holds for every n , so L is a period of f . Boundedness is used exactly once: it supplies one finite common multiple that works for every possible value of $f(n)$.

Topic-based proof models

Section 1 proof model: invariants

Suppose the process is to change a configuration by a legal move and one suspects that a target configuration is unreachable. The first thing to do is not to try all move sequences, since the number of possible sequences grows too quickly. Instead, attach a number, parity, residue class, or weighted sum to the configuration and check how a legal move changes it. If the quantity is unchanged by every legal move, then it is an invariant. If the initial and target configurations have different values of the invariant, the target is impossible.

For a concrete model, color the squares of a board black and white like a chessboard and let

$$I = (\text{number of tokens on black squares}) - (\text{number of tokens on white squares}).$$

If a legal move moves one token from a black square to a white square and another token from a white square to a black square, then the two contributions cancel. The first token changes its contribution by -2 , while the second changes its contribution by $+2$, so the total value of I is unchanged. This calculation is small, but it is the whole proof: every legal move leaves I fixed.

To finish such a problem, one compares the value of I at the beginning and at the claimed end. If the initial arrangement has $I = 0$ and the desired arrangement has all tokens on black squares, giving $I > 0$, then no legal sequence exists. Notice that the conclusion does not say that we failed to find a sequence. It says that every sequence is blocked, because every sequence preserves I . This is the standard logical strength of an invariant argument.

Section 1 proof model: extremal choice

An extremal argument begins by choosing an object that is largest, smallest, leftmost, closest, or otherwise extreme. The purpose is to make all possible competitors one-sided. For instance, if one has a finite set of real numbers and chooses the largest element M , then every other element is at most M . This simple inequality may convert a complicated symmetric condition into a contradiction.

Assume a finite nonempty set S of positive integers has the property that whenever $a, b \in S$, the number $|a - b|$ also belongs to S , unless it is zero. Let m be the least positive element of S . Divide every element of S by m with remainder. If some element $x \in S$ has remainder r with $0 < r < m$, then repeated subtraction of m from x shows that r is also forced by the closure condition. This contradicts the minimality of m . Therefore every element is divisible by m .

This illustrates a common Putnam pattern. The proof is not only that the minimum exists; the proof uses the minimum to show that any smaller positive residue would violate extremality. Whenever a problem contains a finite nonempty set and a closure operation that makes smaller objects, the least element is usually the right one to examine.

Section 2 proof model: ordinary induction

A clean induction proof has two parts: the statement must be strong enough to survive the induction step, and the induction step must explicitly identify where the hypothesis is used. Consider the identity

$$1 + 2 + \cdots + n = \frac{n(n+1)}{2}.$$

The base case $n = 1$ says $1 = 1$, so the formula begins correctly. Assume the formula is true for some n . Then

$$1 + 2 + \cdots + n + (n+1) = \frac{n(n+1)}{2} + (n+1).$$

The induction hypothesis is exactly the replacement of the first n terms by $n(n+1)/2$. Factoring gives

$$\frac{n(n+1)}{2} + (n+1) = \frac{(n+1)(n+2)}{2},$$

which is the same formula with n replaced by $n+1$.

In a Putnam solution, this would be considered elementary, but it shows the discipline needed in harder inductions. One should not write that the result is obvious for $n+1$. One should display the expression at stage $n+1$, isolate the part controlled by the induction hypothesis, and then simplify to the desired form. Harder inductions differ in the algebra, not in the underlying logic.

Section 2 proof model: strong induction

Strong induction is useful when the object of size n naturally breaks into pieces of several smaller sizes. Suppose we want to prove that every integer $n \geq 2$ is a product of primes. The base case $n = 2$ is immediate. Assume now that every integer m with $2 \leq m < n$ is a product of primes. If n itself is prime, we are done.

If n is composite, write $n = ab$ with $2 \leq a, b < n$. By the strong induction hypothesis, both a and b are products of primes, and multiplying those products gives a prime factorization of n .

The point is that ordinary induction would only allow us to use the case $n - 1$, which is irrelevant to a factorization $n = ab$. Strong induction allows all smaller cases, exactly matching the structure of the decomposition. Many Putnam problems have the same feature: a graph splits into connected components, a polynomial factors, an interval splits at a special point, or an integer decomposes into smaller factors. In each case the proof should state clearly that all smaller instances are available.

Section 3 proof model: factorization with a parity condition

Factoring is often a way of discovering the right variables. For example, to solve

$$x^2 - y^2 = 2026$$

in integers, write the left side as $(x - y)(x + y)$. Put $a = x - y$ and $b = x + y$. Then $ab = 2026$, and also a and b must have the same parity because

$$b - a = (x + y) - (x - y) = 2y.$$

Conversely, if a, b have the same parity and $ab = 2026$, then

$$x = \frac{a + b}{2}, \quad y = \frac{b - a}{2}$$

are integers and give a solution.

Since $2026 = 2 \cdot 1013$ has factor pairs with opposite parity, there are no factor pairs (a, b) of the same parity. Therefore the equation has no integer solutions. The important lesson is that the factorization does not merely simplify the equation. It converts the problem into divisors plus a parity condition, and both conditions must be checked.

Section 3 proof model: telescoping

A useful algebraic trick is to create a telescoping expression by adding and subtracting the right intermediate terms. Suppose one wants to show

$$\sum_{k=1}^n \frac{1}{k(k+1)} = \frac{n}{n+1}.$$

The factorization $k(k+1)$ suggests looking for constants A, B with

$$\frac{1}{k(k+1)} = \frac{A}{k} + \frac{B}{k+1}.$$

Multiplying by $k(k+1)$ gives $1 = A(k+1) + Bk$. Comparing coefficients gives $A = 1$ and $B = -1$, so

$$\frac{1}{k(k+1)} = \frac{1}{k} - \frac{1}{k+1}.$$

The sum becomes

$$\left(1 - \frac{1}{2}\right) + \left(\frac{1}{2} - \frac{1}{3}\right) + \cdots + \left(\frac{1}{n} - \frac{1}{n+1}\right),$$

and all intermediate terms cancel. The result is $1 - 1/(n+1) = n/(n+1)$. In contest work, the key is to write enough of the expanded sum that the cancellation pattern is visible and justified.

Section 4 proof model: a nonnegative-square proof

To prove $x + 1/x \geq 2$ for $x > 0$, begin from the nonnegativity of a square:

$$(\sqrt{x} - 1/\sqrt{x})^2 \geq 0.$$

Expanding gives

$$x - 2 + \frac{1}{x} \geq 0,$$

which is exactly $x + 1/x \geq 2$. Equality holds precisely when $\sqrt{x} = 1/\sqrt{x}$, or $x = 1$.

This is the most basic model for equality-case analysis. The proof not only establishes the inequality; it also identifies when equality can occur. In a harder Putnam inequality, the equality case often suggests the normalization or substitution. If equality should occur when all variables are equal, then AM-GM, Cauchy-Schwarz, Jensen, or smoothing may be plausible. If equality occurs at a boundary, then a monotonicity or compactness argument may be better.

Section 4 proof model: Cauchy-Schwarz in Engel form

Cauchy-Schwarz in Engel form says

$$\sum_{i=1}^n \frac{x_i^2}{a_i} \geq \frac{(x_1 + \cdots + x_n)^2}{a_1 + \cdots + a_n}$$

for positive a_i . To prove it, apply the ordinary Cauchy-Schwarz inequality to the vectors

$$\left(\frac{x_1}{\sqrt{a_1}}, \dots, \frac{x_n}{\sqrt{a_n}} \right) \quad \text{and} \quad (\sqrt{a_1}, \dots, \sqrt{a_n}).$$

Their dot product is $x_1 + \cdots + x_n$. Therefore

$$(x_1 + \cdots + x_n)^2 \leq \left(\sum_{i=1}^n \frac{x_i^2}{a_i} \right) \left(\sum_{i=1}^n a_i \right),$$

and division by $a_1 + \cdots + a_n$ gives the desired result.

This proof is worth remembering because it explains why Engel form is natural. The denominators are absorbed by square roots, and the numerator sum appears as a dot product. When a problem contains a sum of fractions and the desired conclusion involves a square of a sum, Engel form should be tested early.

Section 5 proof model: two-variable smoothing

The smoothing principle says that if an expression is symmetric and convex in two variables while their sum is fixed, then replacing those two variables by their average should not decrease the expression. For example, fix $s = x + y$ and study

$$x^2 + y^2.$$

Writing $y = s - x$, we get

$$x^2 + (s - x)^2 = 2 \left(x - \frac{s}{2} \right)^2 + \frac{s^2}{2}.$$

This is minimized at $x = s/2$, so among pairs with fixed sum, the balanced pair minimizes the sum of squares.

In a Putnam inequality, smoothing is rarely announced in the problem. One detects it by finding a symmetric expression and a fixed linear constraint. The proof must then show the two-variable replacement explicitly. Once the expression improves after each averaging step, repeated smoothing pushes the variables toward equality, and the equality case becomes the natural candidate for the extremum.

Section 6 proof model: odd squares modulo 8

To solve a congruence such as

$$x^2 \equiv 1 \pmod{8},$$

check the structure of residues rather than all integers. Every odd integer can be written as $2k + 1$, so

$$x^2 = (2k + 1)^2 = 4k(k + 1) + 1.$$

Since one of k and $k + 1$ is even, the product $k(k + 1)$ is even, so $4k(k + 1)$ is divisible by 8. Hence $x^2 \equiv 1 \pmod{8}$. Even integers have squares congruent to 0 or 4 (mod 8), so the statement is special to odd integers.

The proof is small but the method is powerful. Instead of listing all odd residues, we parameterized the odd integers and used parity inside the expansion. This style scales to many Putnam congruence problems, especially when the modulus is a prime power and one must exploit divisibility inside a product.

Section 6 proof model: multiplicative order

Let a and m be relatively prime. The multiplicative order of a modulo m is the least positive integer d such that $a^d \equiv 1 \pmod{m}$. If $a^n \equiv 1 \pmod{m}$, then $d \mid n$. To see this, divide n by d : write $n = qd + r$ with $0 \leq r < d$. Then

$$1 \equiv a^n = a^{qd+r} \equiv (a^d)^q a^r \equiv a^r \pmod{m}.$$

By minimality of d , the only possible remainder is $r = 0$. Hence $d \mid n$.

The proof shows why order arguments are so efficient. Once a minimal exponent is chosen, every other exponent giving 1 is forced to be a multiple of it. Putnam number theory problems often ask for divisibility of exponents, and this remainder argument is the standard way to extract it.

Section 7 proof model: infinite descent

Infinite descent proves impossibility by showing that any solution would create a smaller positive solution of the same type. For example, suppose positive integers x, y satisfy

$$x^2 = 2y^2.$$

Then x^2 is even, so x is even. Write $x = 2x_1$. Substitution gives

$$4x_1^2 = 2y^2,$$

so $y^2 = 2x_1^2$. Thus y^2 is even, and y is even. Write $y = 2y_1$. Then

$$x_1^2 = 2y_1^2,$$

so (x_1, y_1) is another positive integer solution, but it is smaller than (x, y) .

Repeating the argument would produce an infinite decreasing sequence of positive integers, which is impossible. Therefore no positive integer solution exists. This is the classical irrationality proof for $\sqrt{2}$, but the Putnam lesson is broader: if divisibility forces every variable to share a common factor, divide it out and descend.

Section 8 proof model: divisors as exponent vectors

For the divisor function $d(n)$, write the prime factorization

$$n = p_1^{a_1} p_2^{a_2} \cdots p_r^{a_r}.$$

Every positive divisor of n has the form

$$p_1^{b_1} p_2^{b_2} \cdots p_r^{b_r},$$

where $0 \leq b_i \leq a_i$ for each i . The choices for different primes are independent. There are $a_i + 1$ choices for b_i , so by the multiplication principle

$$d(n) = (a_1 + 1)(a_2 + 1) \cdots (a_r + 1).$$

This proof is more than a formula. It is a bijection between divisors and exponent vectors. When a Putnam problem involves divisors, one should often translate divisors into exponent vectors, because the divisibility relation becomes coordinatewise comparison. That coordinatewise structure is much easier to count or estimate.

Section 8 proof model: Euler's totient formula

To understand $\varphi(n)$, first handle prime powers. Among the integers $1, 2, \dots, p^a$, exactly the multiples of p fail to be relatively prime to p^a . There are p^{a-1} such multiples, so

$$\varphi(p^a) = p^a - p^{a-1} = p^a \left(1 - \frac{1}{p}\right).$$

For general $n = p_1^{a_1} \cdots p_r^{a_r}$, the Chinese remainder theorem shows that being relatively prime to n is equivalent to being relatively prime to each prime power component. Thus the counts multiply:

$$\varphi(n) = n \prod_{p|n} \left(1 - \frac{1}{p}\right).$$

The intuition is that each prime divisor p removes a proportion $1/p$ of the residue classes. The Chinese remainder theorem makes the removals independent across distinct primes. This independence is exactly why multiplicative functions appear so frequently in contest number theory.

Section 9, Exercise 1: subsets counted by size

The identity

$$\sum_{k=0}^n \binom{n}{k} = 2^n$$

has a direct counting proof. The right side counts all subsets of an n -element set, since each element is either chosen or not chosen. The left side counts the same subsets by size: there are $\binom{n}{k}$ subsets of size k , and every subset has exactly one size between 0 and n . Summing over k counts every subset exactly once. This proof is preferable to expanding $(1+1)^n$ when the goal is combinatorial insight. It teaches the standard Putnam habit: identify one set of objects, count it in two ways, and then equate the answers. The partition by subset size is a model for many double-counting arguments, where one side counts globally and the other side counts according to a statistic.

Section 9 proof model: inclusion-exclusion

Inclusion-exclusion for two sets says

$$|A \cup B| = |A| + |B| - |A \cap B|.$$

The reason is that the sum $|A| + |B|$ counts elements in exactly one of the sets once, but counts elements in both sets twice. Subtracting $|A \cap B|$ corrects that overcount. For three sets, the same idea gives

$$|A \cup B \cup C| = |A| + |B| + |C| - |A \cap B| - |A \cap C| - |B \cap C| + |A \cap B \cap C|.$$

Elements in all three sets are counted three times, then subtracted three times, so they must be added back once.

In a Putnam problem, inclusion-exclusion is usually successful when the forbidden conditions are easy to count individually and in intersections. If the intersections become more complicated than the original problem, then a different method may be needed. The proof should always make clear what the underlying universe is and which conditions are being included or excluded.

Section 10 proof model: choosing pigeonholes

The pigeonhole principle says that if N objects are placed into k boxes, then some box contains at least $\lceil N/k \rceil$ objects. The proof is by contradiction. If every box contained at most $\lceil N/k \rceil - 1$ objects, then the total number of objects would be at most

$$k(\lceil N/k \rceil - 1),$$

which is less than N when $\lceil N/k \rceil$ is the least integer not smaller than N/k . This contradicts the existence of N objects.

The practical difficulty is not the proof but the choice of boxes. In number theory the boxes might be residue classes. In geometry they might be small regions covering a larger region. In combinatorics they might be possible values of a statistic. A good solution names the objects, names the boxes, and then states the numerical comparison that forces a collision.

Section 10 proof model: averaging degrees

A typical averaging argument says that if the average of several numbers is A , then at least one number is at least A and at least one is at most A . If all numbers were less than A , their average would be less than A , contradiction. If all were greater than A , their average would be greater than A , contradiction.

For example, in a graph with n vertices and m edges, the sum of degrees is $2m$. Therefore the average degree is $2m/n$. Hence some vertex has degree at least $2m/n$ and some vertex has degree at most $2m/n$. This is often the first step in an extremal graph proof. Choose a vertex with large or small degree, remove it, and apply induction to the remaining graph.

Section 11 proof model: translating a recurrence

For the Fibonacci numbers $F_0 = 0$, $F_1 = 1$, and $F_{n+2} = F_{n+1} + F_n$, the generating function

$$F(x) = \sum_{n \geq 0} F_n x^n$$

can be found by translating the recurrence into an equation. Multiply the recurrence by x^{n+2} and sum over $n \geq 0$:

$$\sum_{n \geq 0} F_{n+2} x^{n+2} = \sum_{n \geq 0} F_{n+1} x^{n+2} + \sum_{n \geq 0} F_n x^{n+2}.$$

The left side is $F(x) - x$, the first right side is $x(F(x))$, and the second is $x^2 F(x)$. Hence

$$F(x) - x = xF(x) + x^2 F(x),$$

so

$$F(x) = \frac{x}{1 - x - x^2}.$$

This derivation is a model for recurrence problems. Initial terms appear because shifting the index leaves out the first few coefficients. Once the recurrence is converted into an algebraic equation, coefficient extraction or partial fractions can finish the problem.

Section 12 proof model: deleting a leaf

A tree with n vertices has $n - 1$ edges. Prove this by induction on n . For $n = 1$, there are no edges. For $n > 1$, every finite tree has a leaf. Remove a leaf v and its unique incident edge. The remaining graph is still connected and has no cycle, so it is a tree with $n - 1$ vertices. By induction it has $n - 2$ edges. Restoring v restores exactly one edge, so the original tree has $n - 1$ edges.

The proof reveals two standard facts about trees: leaves exist, and removing a leaf preserves the tree structure. Putnam graph problems often use this kind of reduction. Rather than trying to understand the whole graph at once, remove a vertex whose local structure is simple and then apply induction to the smaller graph.

Section 13 proof model: power of a point

Power of a point states that if two secants from a point P meet a circle at A, B and C, D , then

$$PA \cdot PB = PC \cdot PD.$$

One proof uses similar triangles. Draw the cyclic quadrilateral with vertices A, B, C, D . The angles subtended by the same chord are equal, so the relevant triangles formed by the two secants are similar. Matching corresponding sides gives the product relation.

The theorem is useful because it converts a geometric configuration into an algebraic equation of lengths. In a contest solution, one should identify the point P , the two lines through it, and the four intersection points with the circle. Once these are named clearly, the product equation follows and can be combined with other metric information.

Section 14 proof model: squared distances and dot products

In vector geometry, the identity

$$\|u - v\|^2 = \|u\|^2 + \|v\|^2 - 2u \cdot v$$

is the algebraic form of the law of cosines. It follows by expanding the dot product:

$$\|u - v\|^2 = (u - v) \cdot (u - v) = u \cdot u - 2u \cdot v + v \cdot v.$$

Thus distance questions can often be converted into dot-product questions.

For example, the condition that two vectors are perpendicular is $u \cdot v = 0$. The condition that two points are at equal distance from a point can be written by equating squared norms. Squared norms avoid square roots and often simplify to linear equations after cancellation. This is why coordinate and vector solutions are frequently shorter than synthetic ones in analytic Putnam geometry.

Section 15 proof model: eigenvalues and the characteristic polynomial

If A is an $n \times n$ matrix and $Av = \lambda v$ for some nonzero vector v , then v is an eigenvector and λ is an eigenvalue. One way to find possible λ is to rewrite the equation as

$$(A - \lambda I)v = 0.$$

Since $v \neq 0$, the matrix $A - \lambda I$ must be singular, so

$$\det(A - \lambda I) = 0.$$

This determinant is the characteristic polynomial.

The important proof detail is the nonzero vector. A homogeneous linear system has a nonzero solution exactly when its coefficient matrix is singular. Thus eigenvalue problems are determinant problems in disguise. In Putnam settings, this allows one to use trace, determinant, rank, and polynomial identities to control linear transformations.

Section 16 proof model: monotone bounded convergence

To prove convergence of a monotone bounded sequence, suppose (a_n) is increasing and bounded above. Let $L = \sup\{a_n : n \geq 1\}$. We claim $a_n \rightarrow L$. Given $\varepsilon > 0$, the number $L - \varepsilon$ is not an upper bound, so some term a_N satisfies $a_N > L - \varepsilon$. Since the sequence is increasing, for every $n \geq N$ we have

$$L - \varepsilon < a_N \leq a_n \leq L.$$

Therefore $|a_n - L| < \varepsilon$ for all $n \geq N$.

This proof is often the hidden engine behind Putnam limit problems. One first proves monotonicity, then proves a bound, and then uses the completeness of the real numbers through the supremum property. The final identification of the limit may require passing to the limit in a recurrence or equation.

Section 17 proof model: the intermediate value theorem

The intermediate value theorem is useful when an equation is hard to solve exactly but signs are easy to determine. If f is continuous on $[a, b]$, $f(a) < 0$, and $f(b) > 0$, then there is $c \in (a, b)$ with $f(c) = 0$. The intuition is that a continuous graph cannot move from negative height to positive height without crossing the axis.

In a proof, define the correct function first. For example, to show that $x^3 + x - 1 = 0$ has a root in $(0, 1)$, put $f(x) = x^3 + x - 1$. Then $f(0) = -1$ and $f(1) = 1$, and f is continuous because it is a polynomial. The theorem gives the root. If uniqueness is also required, show that $f'(x) = 3x^2 + 1 > 0$, so f is strictly increasing and cannot cross the axis twice.

Section 18 proof model: the mean value theorem

The mean value theorem says that if f is continuous on $[a, b]$ and differentiable on (a, b) , then

$$f(b) - f(a) = f'(c)(b - a)$$

for some $c \in (a, b)$. This turns a difference of values into a derivative estimate. If $|f'(x)| \leq M$ on the interval, then

$$|f(b) - f(a)| \leq M|b - a|.$$

This is a common way to prove inequalities. For example, for $x > 0$, apply the theorem to $f(t) = \log t$ on the interval between 1 and x . Since $f'(t) = 1/t$, derivative bounds give estimates for $\log x$. The proof

should always specify the interval and verify differentiability, since the theorem is a compactness statement in disguise.

Section 19 proof model: interpolation uniqueness

Polynomial interpolation says that a polynomial of degree at most n is determined by its values at $n + 1$ distinct points. If two such polynomials P and Q agree at $n + 1$ distinct points, then $P - Q$ has at least $n + 1$ roots. But $P - Q$ has degree at most n , so it must be the zero polynomial. Hence $P = Q$.

The proof is short, but the idea is powerful. To prove a polynomial identity, it can be easier to check sufficiently many values than to expand both sides. Putnam polynomial problems often hide this method. Once a degree bound is known, roots become a counting resource, and too many roots force the polynomial to vanish identically.

Section 20, Exercise 2: bounded additive functions

A functional equation solution should begin by substituting special values that simplify the equation. If the equation contains $f(x + y)$, test $y = 0$. If it contains products, test $x = 0$, $x = 1$, or values that make factors vanish. The goal is to discover fixed values such as $f(0)$, injectivity, additivity, or multiplicativity. For instance, suppose $f : \mathbb{R} \rightarrow \mathbb{R}$ satisfies

$$f(x + y) = f(x) + f(y)$$

and is bounded on some interval of positive length. First, putting $y = 0$ gives $f(0) = 0$. Then putting $y = -x$ gives $f(-x) = -f(x)$. Repeated addition gives $f(nx) = nf(x)$ for integers n , and then $f(qx) = qf(x)$ for rational q . The boundedness condition rules out the wild additive solutions and forces continuity at 0. Once continuity is known, rational approximation gives $f(x) = cx$ for all real x .

The proof shows how strategy synthesis works. Algebraic substitutions give rational linearity, while an analytic condition supplies continuity. Many advanced Putnam problems require exactly this mixing of tools.