Some Thoughts on Writing for the Putnam

Bruce Reznick[14]

University of Illinois at Urbana-Champaign

What is the Putnam?

This chapter describes the process of composing problems for the William Lowell Putnam Mathematical Competition. Inevitably, this leads to the more general issue of mathematical problem writing. I shall be anecdotal and probably idiosyncratic, and do not purport that my opinions are definitive or comprehensive. The reader should not look for "How to Pose It," but rather sit back and enjoy the heart-warming tale of a boy and his problems.

The following is a quotation from the official brochure (Putnam, 1992). The William Lowell Putnam Mathematical Competition "began in 1939 and is designed to stimulate a healthful rivalry in mathematical studies in the colleges and universities of the United States and Canada. It exists because Mr. William Lowell Putnam had a profound conviction in the value of organized team competition in regular college studies." The Putnam, as it is universally called, is administered by the Mathematical Association of America. It is offered annually (since 1962, on the first Saturday in December) to students who have not yet received a college degree. In 1991, 2,375 students at 383 colleges and universities took the exam. (For more on a history of the Putnam, see [**PutnamI, PutnamII**].) The problems and solutions for Competitions through 1984 are in [**PutnamI, PutnamII**]. The problems, solutions, and winners' names are also published in the *American Mathematical Monthly*, usually about a year after the exam.

The Putnam consists of two independent 3-hour sessions, each consisting of six problems arranged roughly in order of increasing difficulty. The exam is administered by proctors who cannot comment on its content. Contestants work alone and without notes, books, calculators, or other external resources. They are ranked by their scores, except that the top five are officially reported en bloc. Teams are preselected by their coach, and team rankings are determined by the sum of the ranks (not the sum of the scores).

> The examination will be constructed to test originality as well as technical competence. It is expected that the contestant will

[14]This chapter originally appeared in [**Re4**]. The author was supported in part by the National Science Foundation.

> be familiar with the formal theories embodied in undergraduate mathematics. It is assumed that such training, designed for mathematics and physical science majors, will include somewhat more sophisticated mathematical concepts than is the case in minimal courses.
>
> Questions will be included that cut across the bounds of various disciplines and self-contained questions which do not fit into any of the usual categories may be included. It will be assumed that the contestant has acquired a familiarity with the body of mathematical lore commonly discussed in mathematics clubs or in courses with such titles as "Survey of the foundations of mathematics." It is also expected that self-contained questions involving elementary concepts from group theory, set theory, graph theory, lattice theory, number theory, and cardinal arithmetic will not be entirely foreign to the contestant's experience. (from official brochure)

Between 1969 and 1985, I participated in all but three Putnams. I competed four times, the last two wearing the silks of the Caltech team, which placed first. As a graduate student and faculty member, I coached or assisted at Stanford, Duke, Berkeley, and Urbana. I was a grader in 1982 and a member of the Problems Subcommittee for the 1983, 1984, and 1985 Competitions. (I had been living quietly in Putnam retirement when Alan Schoenfeld invited me to this conference.[15])

Who writes the Putnam?

The Problems Subcommittee of the MAA Putnam Committee consists of three question writers, who serve staggered 3-year terms. The most senior member chairs the Subcommittee. During my service, the Problems Subcommittee consisted of successive blocks of three consecutive people from the following list: Doug Hensley, Mel Hochster, myself, Richard Stanley, and Harold Stark. Outgoing members are invited to suggest possible replacements, but these are not acted on immediately. This is an old-boy network, but in practice one often suggests the names of strangers whose problems one has admired. (It turned out that the 1985 group inadvertently consisted of three Caltech alumni, and, as chair, I was relieved when the Harvard team won.)

The rest of the Putnam Committee consisted of three permanent members, two of whom (Jerry Alexanderson and Leonard Klosinski) arranged the massive logistics of the Competition and a liaison with the Problems Subcommittee (Abe Hillman, then Loren Larson). They do an incredible amount of work, which is not germane to this essay, but should not go unappreciated. They also provide a valuable final check to the Subcommittee's work.

How was the 1985 Putnam Written?

The following is an overview of how the 1985 Putnam was written. In November 1984, I wrote a welcoming instructional letter to Richard and Harold, describing

[15]This article was originally written for the conference "Mathematical Thinking and Problem Solving", held in Berkeley on December 14–16, 1989.

our timetable and goals. In early December, each of us circulated about a dozen problems to the other two. After a decent interval, we circulated solutions to our own problems and comments about the others and added some more into the pot. A few more rounds of correspondence ensued. In March, we met for a weekend with the rest of the Committee to construct the Competition. After a small flurry of additional correspondence, the material was handed over to the logistics team.

The greeting letter of the chair of the Problems Subcommittee is a quilt to which each chair adds (or rips up) patches. The following excerpts, lightly rewritten for style, are thus simultaneously traditional and fully my own responsibility.

> It used to be said that a Broadway musical was a success if the audience left the theater whistling the tunes. I want to see contestants leave the Putnam whistling the problems. They should be vivid and striking enough to be shared with roommates and teachers.
>
> Security should be a major concern ... the problems should be handwritten or typed by ourselves, and our files should either be unmarked or kept home. Putnam problems ought to be pretty enough that you want to tell your friends about them. Do your best to resist this temptation.
>
> Lemmas in research papers are fair game, but material from well-known textbooks or problem collections are not. (If you submit a problem from a known source, please include this information.) Seminar material is OK unless undergraduates were present, and anything you have taught to an undergraduate honors class ought to dry out for a few years before gaining eligibility. In general, problems from other people are not reliably secure unless your source can vouch for their originality.
>
> As for the problems themselves, my feeling is that any problem solved by only one or two contestants is a failure, no matter how beautiful it might be. In the last couple of years, we have sought to turn A-1 and B-1 into "hello, welcome to the exam" problems, and their relative tractability has been appreciated. It is better to require one major insight than several minor ones (partial credit is undesirable). It is better to write a streamlined problem without many cases, so that we test perceptiveness, rather than stamina. Proofs by contradiction are, in general, unsuitable, both because they are ugly, and because they are harder to grade.
>
> Although concern for the graders is not our primary consideration, we should keep them in mind. There is no reason to exclude a problem such as 1983 B-2, merely because there are many different legitimate proofs. On the other hand, we must be at pains to write unambiguous questions even at the expense of simplicity in the phrasing. Answers in a particular numerical form are often desirable so that students won't puzzle over the phrase "simplest terms"; this is one reason that the current year stands in for "n". Answers that turn out to be 0, 1, π,

> $\sqrt{2}$, etc., should be avoided to eliminate the lucky guess, and we should not present problems in which the solution is easy to guess but hard to prove. (The reverse is preferable.) I confess to a predilection for "garden variety" mathematical objects, such as powers of 2, binomial coefficients, pentagons, sines, cosines, and so on. I dislike problems with an elaborate notation, whose unraveling is a major portion of the solution.

What does the Putnam mean?

The previous three sections have given a theoretical description of how the Putnam works, and what it is intended to accomplish. I turn now to the Putnam in reality.

There is some evidence that the Putnam achieves its intended goals. Many schools run training sessions for contestants, in which interesting mathematics and useful techniques of problem solving are presented. A successful individual performance on the Putnam leads to fame and glory and an increased probability of a fellowship. (However, the results are announced in March, very late for seniors applying to graduate school.) There is also money; I was entertained for years by the William Lowell Putnam Stereo System. More importantly, a contestant can properly be satisfied in solving any Putnam problem, though this is tempered by the (larger number of) problems one does not solve. Putnam problems have occasionally led to research, and a problem may stick in a contestant's mind for years. The ultimate source of [**Re2**] was 1971 A-1.

The phrase "Putnam problem" has achieved a certain cachet among those mathematicians of the problem-solving temperament and is applied to suitably attractive problems which never appeared on the exam. One motivation for my joining the Problems Subcommittee was the aesthetic challenge of presenting to the mathematical community a worthy set of problems. In fact, the opportunity to maintain this "brand name of quality" was more enticing to me than the mere continuation of an undergraduate competition. Of course, the primary audience for the Putnam must always be the students, not one's colleagues.

At the same time, the Putnam causes a few negative effects, mainly because of its difficulty. Math contests are supposed to be hard, and the Putnam is the hardest one of all. In 1972, I scored less than 50% and finished seventh. In most years, the median Putnam paper has fewer than two largely correct solutions. For this reason, the first problem in each session is designed to require an "insightlet", though not a totally trivial one. We on the Committee tried to keep in mind that median Putnam contestants, willing to devote one of the last Saturdays before final exams to a math test, are likely to receive an advanced degree in the sciences. It is counterproductive on many levels to leave them feeling like total idiots.

Success on the Competition requires mathematical ability and problem-solving experience, but these are not sufficient; a "Putnam" temperament is also necessary. A contestant must be able to work quickly, independently, and without references and be willing to consider problems out of context. I have been saddened by reports of students who were discouraged in their academic careers by a poor performance

on the Putnam. Fortunately for the mathematical community, there are many excellent, influential, and successful mathematicians who also did badly on the Putnam. As a result, the absence of Putnam kudos has a negligible effect on one's career. (At the same time, I confess to enjoying the squirmy defensiveness that the term "Putnam" evokes in some otherwise arrogant colleagues of the "wrong" temperament. They loudly deny an importance to the Competition that nobody else asserts.)

Among those who do very well on Putnam problems, there is little hard evidence that doing extremely well is significant. The best papers usually average about twice as many correct solutions as the thirtieth. My impression is that the likely future mathematical outputs of the writers are comparable.

In sum, the Putnam plays a valuable, but ultimately inessential, role in undergraduate mathematics. This is a test; this is only a test.

What makes a good Putnam problem?

Other considerations besides pure problem aesthetics afflict the Putnam writers. We wish to have a balance of questions in various subject areas and solving styles. We need an "easy" question for A-1 and B-1. As college math teachers, we are often astonished at how poorly we know what it is that our own students do and do not understand. This is magnified on the Putnam, in which contestants come from hundreds of different programs. The Committee tries to be sensible. It's unreasonable to have the trace of a matrix in one of the easier problems, but we used it in 1985 B-6, on the grounds that a contestant who had not heard of a trace would probably be unable to do the problem anyway.

We want to test, if possible, abstract problem-solving ability, rather than classroom knowledge; maturity "yes," facts "no." We try to avoid the traditional corpus of problem-solving courses to minimize the reward in "studying" for the Putnam. This leads to a tradeoff between familiarity and quality. We occasionally receive a complaint that a problem is not new or has even appeared in material used to train Putnam competitors at a particular school. This is unfortunate, but probably inevitable. It would be easy to write an exam with twelve highly convoluted, certifiably original, and thoroughly uninteresting problems.

Otherwise acceptable questions have been rejected on the grounds that they "sound" familiar or "must have appeared somewhere," even when no member of the Committee can cite a reference. Here's an example: A projectile is to be fired up a hill which makes an angle α with the horizontal. At what angle should the projectile be launched in order to maximize the distance it travels?

One April, the day after the Committee completed its deliberations, someone discovered that our A-1 was a problem posed in the most recent *Two-Year College Mathematics Journal*. Fortunately, we had bequeathed an easy problem to the following year's exam, and a few phone calls resolved the crisis. I do not know what we would have done if this had happened in November.

What follows are some representative comments evoked by the first round of proposed problems in the Subcommittee:

"Seems routine, too easy."
"I found the computations too messy, and it was easy to head off in the wrong direction."
"Can't use, it's been around for years."
"Hard (or did I miss something). A good problem."
"The trouble with this one is that an intuitive guess ... is incorrect. OK, not inspired."
"I was stumped, but it's a nice hard problem."
"I couldn't do this one either."

Every Putnam I helped write contained at least one problem I could not do on first sight. More comments follow:

"This looks very messy. I saw nothing that motivated me to take up pencil and scribble."
"I wasn't lit up by waves of excitement."
"We must use this one, I love it."
"Yes, yes, yes."

(The last block suggests the sensuality of a good problem to the discriminating solver. For more on this subject, see [**Re1**].)

How are Putnam problems polished?

The Committee acts by consensus; I do not ever recall voting on a problem. Most problems have one primary author, although the full group polishes the final version. Often, though, this version is a special case of the original problem. We tried to have at least one problem from calculus, geometry, and number theory on the exam. I was always amazed at the ability of the Committee to find unfamiliar problems in such fully excavated fields. It is hard to write serious algebra problems that are not basically manipulative, rather than conceptual, because we can assume so little knowledge. In analysis, any integrals must be innocent of measure theory.

We do not accept a problem until we have seen a solution written out in full; sometimes we produce more than one solution. It is not unusual for contestants to find new (and better) solutions. One problem (1983 B-2) evolved from the remarkable fact, familiar to our silicon friends, that a nonnegative integer n has a unique binary representation. Let $f(n)$ denote the number of ways that n can be written in the form $\sum a_i 2^i$, if the a_i's are allowed to take the values 0, 1, 2, or 3. It turns out that $f(n) = \lfloor n/2 \rfloor + 1$. (Here, and below, $\lfloor x \rfloor$ represents the largest integer $\leq x$.) This problem has at least three different solutions:

(1) by generating functions — $\sum f(n)x^n = \{(1-x)(1-x^2)\}^{-1}$,
(2) by induction — use the recurrence $f(2n) = f(2n+1) = f(n) + f(n-1)$, or
(3) by direct manipulation — writing $a_i = 2b_i + c_i$, where b_i and c_i are 0 or 1, gives a bijection onto sums $n = 2k+m$, where $k = \sum b_i 2^i$ and $m = \sum c_i 2^i$.

I have explored this topic more extensively elsewhere [**Re3**].

Several times, there was true collaboration. Doug Hensley called me to say that he wanted a problem in which an algorithm terminated because a certain nonnegative integral parameter decremented by 1 after each iteration. This reminded me of a situation I was playing with, in which n was replaced by $n + \lfloor\sqrt{n}\rfloor$; the relevant parameter is $(n - \lfloor\sqrt{n}\rfloor^2) \pmod{\lfloor\sqrt{n}\rfloor}$, which decrements by 1, usually after two iterations. It follows that one eventually reaches a perfect square (see 1983 B-4).

Another time, Mel Hochster had been playing with a problem using the tips of the hands of an accurate clock (as we ultimately phrased it — we received complaints from students who were only familiar with digital clocks!). This was a nice "trapdoor" situation. A reasonably competent student could parametrize the positions of the tips, and after a half-hour of calculation, solve the question. The intended insight was to make a rotating set of coordinates in which the minute hand is fixed, so only the hour hand is rotating. We would have a problem, if only we could find the right question. It occurred to me to look at the distance between the tips when that distance was changing most rapidly. Solving it the long way, I uninspiredly computed an answer that shouted out, "You idiot, use the Pythagorean theorem!" In fact, the derivative of the difference vector from one tip to the other has constant magnitude and is normal to the hour vector, and the distance is changing most rapidly when the difference vector and its derivative are parallel (see 1983 A-2).

The length of service as a Putnam writer seems optimal. In my first year, I was bursting with problems I had saved for the exam and discovered that some were unsuitable. In the second year, I tried to rework the leftovers and develop some techniques for consciously writing other problems (these will be discussed below). (A few Putnam rejects have appeared in the *Monthly* problem section, where the lack of time constraints relieves concern over messy or evasive algebra.) By the third year, I felt drained of inspiration; in fact, my impression all three years was that the chair placed the fewest problems on the Putnam. Service on the Subcommittee was also beginning to have an adverse effect on my research. Ordinarily, a mathematician tries to nurture a neat idea in hopes that it will grow into a theorem or a paper. I found that I was trying to prune my ideas so that they would fit on the exam. Bonsai mathematics may be hazardous to your professional health!

Did Archimedes use δ's and ϵ's?

The story of one of my favorite problems (1984 B-6) serves as an object lesson in theft. In the course of researching the history of the Stern sequence and Minkowski's ?-function, I had run across a beautiful example of Georges de Rham. Let P_0 be a polygon with n sides, trisect each side, and snip off the corners, creating P_1, a polygon with $2n$ sides. Iterate. The boundary of the limiting figure, P_∞, has many interesting counterintuitive properties (see the last paragraph of the discussion of this problem in [**PutnamII**]). For example, it is a smooth convex curve which is flat almost everywhere. I was rather pleased with myself for having noticed a property of P_∞ itself. Suppose one corner, snipped from P_i, has area A. Then each of the two new adjacent corners snipped from P_{i+1} has one-third the altitude and one-third the base of the original corner, and so has area $A/9$. Further, if P_0 is a triangle, then P_1 is a hexagon whose area is two-thirds the area of P_0. This

information can be combined with the formula for the sum of a geometric series with $r = 2/9$ to show that the area of P_∞ is four-sevenths the area of P_0. A sneaky new problem that requires only precalculus — Putnam heaven! I had mentioned this result in a seminar a few years before, but no undergraduates attended, and I was confident of security.

The Tuesday before the competition, I attended a Pi Mu Epsilon lecture on the approaches of Archimedes to calculus, given by Igor Kluvanik, an Australian mathematician visiting Urbana. To my horror, I learned that Archimedes had stolen my method in order to compute the area under a parabola! At least one colleague noticed that I lost my color, and I told her that I could explain the circumstances in about a week. Fortunately, our students did not do unusually well on that problem.

By the way, we stated this problem so that P_0 was an equilateral triangle of side 1, and asked for the area in the form $a\sqrt{c}$, where a is rational, and c is integral. The majority of solvers assumed, incorrectly, that P_∞ had to be the circle inscribed in the triangle and derived an answer involving π. (When I mentioned this problem at a colloquium, a famous mathematical physicist in the audience audibly made the same guess.) The reader might find it amusing to consider the following variation, in which the resulting figure is not a circle, but is piecewise algebraic: Suppose that, rather than a trisection, each side is split in ratio $1 : 2 : 1$ before the corners are snipped off. Describe P_∞.

What makes a good Putnam problem?

The instructions sent to composers do not include a description of the characteristics of a good Putnam problem. The aesthetic seems to be fairly universal among dedicated problem solvers and can be applied more generally to describing good mathematical problems. (Perhaps this reluctance to be specific also reflects the mathematician's cowboy taciturnity on such woolly subjects.)

I will hazard some definitions. A mathematical problem is simply a mathematical situation in which some information is implied by other information. The principal characteristics of a good problem are simplicity, surprise, and inevitability. By *inevitability*, I mean two things: once you've solved the problem, you cannot look at it without also seeing its solution; once you see the problem, you feel you *must* solve it. (A tacit rite of passage for the mathematician is the first sleepless night caused by an unsolved problem.)

As an undergraduate competitor, I told my housemates that doing well on the Putnam reflected one's ability to do very quickly, other people's tricky, solved problems. I'll stand by that today. Three of the most important preliminary questions a problem solver must face are: (a) Is there a solution? (b) What do I need to know to find the solution? and (c) What does the solution look like? These questions are all answered in advance for the Putnam competitor. You know that there *is* a solution, which is probably short and clever and does not require a great deal of knowledge. You know that you will recognize the solution when you see it. Tables and computer data and other references are irrelevant, and inaccessible in any case. You cannot collaborate, or even describe the problem to someone else in hopes of understanding it in the retelling. It is for these reasons that I am extremely

unhappy when I hear that some problem-solving courses use the Putnam as a final exam. Chocolate decadence cake à la mode is a delicious dessert, but makes an unfilling main course.

This discussion begs the larger question of the role of problems within mathematics. Simplicity and surprise may be enjoyable, but they do not accurately characterize much of the mathematician's world, in which correct insights are often wrested after much reflection from a rich contextual matrix and are as snappy as a tension headache.

Okay, so how do you sit down and create a Putnam problem?

Okay, so how do you sit down and create a Putnam problem? One way is to keep your eyes open for anything in your own work that looks like a Putnam problem. You can make a votive offering to Ineedalemma, the tutelary goddess of mathematical inspiration.

You can also be somewhat more systematic. I was fortunate to have a father who addressed very similar questions in his own work: comedy writing. He wrote a book about writing jokes, from which I take the following quotation [**ReS**, p. 15].

> Very few writers can pound out a huge batch of jokes week after week relying solely on sheer inspiration. They need the help of some mechanical process. When a writer "has to be funny by Tuesday" he's not going to wait for hot flashes of genius, especially if he doesn't happen to be feeling too "geniusy." ... Switching is the gag writer's alchemy by which he takes the essence of old jokes from old settings and dresses them up in new clothes so they appear fresh.

(Putnam punks such as myself are often inveterate punsters. Punning requires the rapid formal combinatorial manipulation of strings of symbols, without much concern for content. This skill is also very helpful on the Putnam. More serious connections between humor and creativity are discussed in [**Koe**].)

The details of switching problems and switching jokes are substantially different, but the principle is the same. The following is one practical illustration of problem switching. I heard a seminar speaker refer to a beautiful result of Mills: There is a positive number α with the property that $p_n = \lfloor \alpha^{3^n} \rfloor$ is a prime for every $n \geq 1$. The construction is recursive, based on the observation that α must lie between $(p_n)^{3^{-n}}$ and $(p_n + 1)^{3^{-n}}$, and there is always a prime between any two consecutive cubes. It occurred to me that something similar might be wrought out of the simpler expression $\lfloor \alpha^n \rfloor$. One of the most familiar properties of α^n is that it is always even if α is an even integer (and $n \geq 1$) and always odd if α is odd. The most counterintuitive behavior for $\lfloor \alpha^n \rfloor$ would thus be for it to alternate between even and odd. If you start with $\lfloor \alpha \rfloor \geq 3$, then Mills's interval argument will work, and this is how 1983 A-5 came to life.

Although this problem was solved by fewer contestants than we had hoped, perhaps some of them later realized that the alternation of even and odd is basically irrelevant to the problem, and that any pattern of parities $\pmod 2$ can be achieved

using the same proof, as well as any pattern (mod m). Later on, William Waterhouse found an explicit algebraic integer with the property that $\lfloor \alpha^n \rfloor$ alternates in parity and submitted this version of A-5 to the *Monthly* Problems Section.[16] I refereed it, and we received dozens of correct solutions; see [**AEMS, BDSW1**].

Another way to create Putnam problems is via Fowler's method. Gene Fowler once explained that it's very easy to write. All you have to do is sit at a typewriter and stare at a sheet of blank paper until blood begins to appear on your forehead. I often applied this technique at less exciting seminars and colloquia, when my neighbors though I was doodling. I'd take a combination of simple mathematical objects and stare at them until I could see a Putnam problem. Sometimes it worked. For example, 1984 A-4 asks for the maximal possible area of a pentagon inscribed in a unit circle with the property that two of its chords intersect at right angles.

Contrary to popular opinion, it's unhelpful to read through old problem books very much for inspiration, because subconscious plagiarism is a great danger, and our larger audience is very alert. (It might be more useful to look through the back of books, because switches based on solutions are less transparent.)

So how do you sit down and create a Putnam problem? Let's apply Pólya's rules and generalize the question. How do you sit down and create? This is a very difficult and personal question. (It might not even have an answer; our romantic culture tends to identify the results of algorithmic thinking as mechanical, rather than creative.)

In the end, you can do everything you can, rely on the rest of the Committee for inspiration, and visualize two thousand fresh minds on the first Saturday in December, eager to be challenged.

[16]Editors' note: one possible α is the positive root of the quadratic $x^2 - ax - b$, where a is odd, b is even, and $b < a$.