THE INVARIANT m_0 OF SAWIN–SHUSTERMAN–STOLL FOR THE PAIRS $(1+kx,\,l)$: A STOPPING RULE, EVERY PRIME l, AND A PROVED RANGE

KOREA SUPERINTELLIGENCE LABS

Abstract. Sawin, Shusterman and Stoll attach to a pair (c,d) of integer polynomials with $c(0),d(0)\neq 0$ an integer $m_0(c,d)$: the least M such that for every $m>M$, every pair (a,b) with $\deg a,\deg b<m$, $ab\equiv cd \bmod x^m$ and $\|a\|+\|b\|\le\|c\|+\|d\|$ already satisfies $ab=cd$, where $\|\cdot\|$ is the squared Euclidean length of the coefficient vector. It controls the threshold beyond which the non-reciprocal part of $x^Nc(x^{-1})+d(x)$ is irreducible. For the two-parameter family $(c,d)=(1+kx,\,l)$ they prove $m_0=1$ on a range, and remark that outside it $m_0(1+kx,l)\le 2$ *appears* to hold, with the two exceptions $m_0(1+2x,4)=4$ and $m_0(1+3x,9)=3$; a proof "should be possible", they write, "but the arguments seem to get rather technical". No verification range is stated. We contribute three things. First, a stabilisation lemma: for a weakly robust pair with $(cd)(0)\neq 0$, the property "every element of T_m satisfies $ab=cd$" is monotone in m. This is what makes the authors' own procedure — compute $T_1,T_2,\dots$ until $ab=cd$ throughout — a stopping rule, and it turns the infinite assertion $m_0(c,d)\le M$ into a single condition on the one set T_{M+1}; the source does not state it. Second, an infinite unconditional sub-family: $m_0(1+kx,l)\le 2$ for *every* $k\ge 1$ and *every* prime l, with no computation, by a sum-of-squares certificate; the source's lemma reaches prime l only for $l\ge 5$ and $k\le l^2/2$. Third, two narrowing inequalities which confine a hypothetical further exception to the band $2\le a_0\le b_0\le k\le\sqrt{a_0^2+b_0^2}$, where $a_0b_0=l$. They leave k in an interval of length about $a_0^2/2b_0$ rather than in a box, and enumerating that interval — each of the two levels of the enumeration being a line meeting a disk, which is what makes it affordable — proves the remark's bound $m_0\le 2$ for *every* $k\ge 1$ and every l with $2\le l\le 4\cdot 10^4$ such that $(1+kx,l)$ is weakly robust and (k,l) is not one of the two exceptions; no upper bound on k is imposed. We also certify the two exceptional values themselves, $m_0(1+2x,4)=4$ and $m_0(1+3x,9)=3$, and that the constant 2 cannot be lowered. All of this is machine-checked in Lean 4.

1. Introduction

The objects. Fix $c,d\in\mathbb{Z}[x]$ with $c(0),d(0)\neq 0$ and consider the polynomials *with a large gap*

$$f_N(x)=x^Nc(x^{-1})+d(x),\qquad N>\deg c+\deg d.$$

Sawin, Shusterman and Stoll [6] give, for pairs (c,d) satisfying $\gcd(\tilde c,d)=1$, $c\neq\pm d$ and a robustness condition recalled below, an effective bound $N_0=N_0(c,d)$ such that the non-reciprocal part of f_N is irreducible over $\mathbb{Q}$ for all $N>N_0$. Their N_0 is much smaller than the earlier bounds of Schinzel [4] and of Filaseta–Ford–Konyagin [1], at the cost of the robustness hypothesis, which those results do not need; and they describe an algorithm computing such an N_0 for a given pair. Both the bound and the algorithm run through one integer invariant of the pair.

Write $\|f\|$ for the *weight* of $f\in\mathbb{Z}[x]$, "the squared Euclidean length of its coefficient vector (i.e., the sum of the squares of the coefficients)" [6, §1], and $f|_m$ for f truncated to degree $<m$, that is, the remainder of f on division by x^m. For $m\ge 0$ put

$$T_m(c,d)=\big\{(a,b):a,b\in\mathbb{Z}[x],\ \deg a,\deg b<m,\ ab\equiv cd \bmod x^m,\ \|a\|+\|b\|\le\|c\|+\|d\|\big\},\tag{1}$$

the set displayed as (4.1) of [6]. Call m *stable* for (c,d) when every $(a,b)\in T_m(c,d)$ satisfies $ab=cd$. Corollary 4.4 of [6] says that if (c,d) is weakly robust in the sense recalled in Section 2, then all m beyond $(1+\deg c+\deg d)2^{\|c\|+\|d\|-1}$ are stable; and $m_0(c,d)$ denotes the optimal such bound, so

that

$$m_0(c,d) \le M \quad \Longleftrightarrow \quad \text{every } m > M \text{ is stable for } (c,d).$$

The threshold fed to the irreducibility theorem is $N_0 = 2\max\{\deg c, \deg d, m_0\}$, or $N_0 = \max\{4\deg(cd), 2m_0\}$ in the remaining case, both read off the proof of [6, Lemma 4.5]. So the size of m_0 is the practical content of the method, and [6, §5], "Growth of m_0", is devoted to replacing the exponential a-priori bound by something small on concrete families.

The question. The smallest interesting family is $(c,d) = (1+kx, l)$ with $l \neq \pm 1$ and $k \neq 0$; changing the sign of x or of d normalises $k, l > 0$, as [6] notes, so we take $k \ge 1$ and $l \ge 2$ throughout. Here $cd = l + klx$ and $\|c\| + \|d\| = 1 + k^2 + l^2$. Lemma 5.2 of [6] proves

$$m_0(1+kx,\, l) = 1 \quad \text{for } l \ge 5,\ k \le l^2/2, \text{ and } \begin{cases} k \ge 1 & \text{if } l \text{ is prime,} \\ k > \sqrt{p^2 + l^2/p^2} & \text{otherwise,} \end{cases}$$

where p denotes the least prime factor of l. What lies outside that range is the subject of the paper's Remark 5.3, which we quote in full:

> For the robust pairs $(1+kx, l)$ with $l > 1$ that are not covered by Lemma 5.2, it appears that $m_0(1+kx, l) \le 2$, except for $m_0(1+2x, 4) = 4$ and $m_0(1+3x, 9) = 3$. It should be possible to prove this, but the arguments seem to get rather technical.

The remark is stated as an observation and carries no proof, no proof sketch, and — this is the point of departure for the present paper — *no verification range*: nothing in [6] says over which k and l the assertion was tested.

What is settled here.

(1) **A stopping rule becomes one.** Immediately after Corollary 4.4 the authors write that "for any given pair (c,d), we can effectively determine the optimal bound m_0 by successively computing the sets (1) for $m = 1, 2, \ldots$, until $ab = cd$ for all $(a,b) \in T_m$". That is a correct halting criterion precisely when stability is monotone in m: otherwise a later T_m could sprout a fresh pair with $ab \neq cd$ and the loop would stop too early. Theorem 3.1 supplies the monotonicity for every weakly robust pair with $(cd)(0) \neq 0$, uniformly in m and in (c,d). Its corollary is what makes everything below finite: $m_0(c,d) \le M$ is a single condition on the one set T_{M+1}.

(2) **An infinite sub-family of Remark 5.3, proved.** Theorem 5.2: $m_0(1+kx, l) \le 2$ for every $k \ge 1$ and every prime l, with no computation at all. Lemma 5.2 of [6] reaches prime l only for $l \ge 5$ and $k \le l^2/2$, so $l = 2$, $l = 3$ and every $k > l^2/2$ are new. The proof rests on the sum-of-squares certificate $12\beta^2 + 4(a_2+1)^2 + (4\gamma - 1)^2 \le 5$.

(3) **A proved range for the rest.** Lemmas 6.1 and 6.2 confine a hypothetical exception with non-trivial constant-term splitting $l = a_0 b_0$, $2 \le a_0 \le b_0$, to $b_0 \le k \le \sqrt{a_0^2 + b_0^2} \le b_0\sqrt{2}$: a two-parameter band rather than a four-dimensional search. Theorem 7.1 then proves Remark 5.3's bound $m_0 \le 2$ — with the hypothesis weakened from robust to weakly robust, and for *all* such pairs of the box rather than only those Lemma 5.2 misses — for $1 \le k \le 1000$, $2 \le l \le 1000$.

(4) **The exceptions and the constant.** Proposition 10.1 certifies $m_0(1+2x, 4) \ge 3$ and $m_0(1+3x, 9) \ge 3$, so the two exception clauses of Theorem 7.1 are not decoration; Proposition 10.2 certifies $m_0(1+3x, 2) \ge 2$, so the constant 2 cannot be replaced by 1. In both propositions the mathematics is the source's and only the certification is ours.

(5) **Forty times that range, and no bound on k at all.** Theorem 9.1: the remark's bound $m_0 \le 2$ holds for *every* $k \ge 1$ and every l with $2 \le l \le 4 \cdot 10^4$, for every weakly robust $(1+kx, l)$ outside the two exceptions. The k-direction here is not an enlarged search but a deleted hypothesis: by Lemmas 6.1 and 6.2 a counterexample has $b_0 \le k \le \sqrt{a_0^2 + b_0^2}$, so the admissible k form an interval of length about $a_0^2/2b_0$, which can be enumerated instead

of bounded. What makes the l-direction reach forty times as far is the subject of Section 8: both levels of the remaining search are lines meeting disks, and the chord bound replaces two windows of width $O(l)$ by a level-1 window of width $O(a_0^3/b_0)$ and a level-2 test that is empty at all but a few hundred of the tens of millions of candidates reaching it.

(6) **The two exceptional values, exactly.** Proposition 10.3: $m_0(1+2x,4)=4$ and $m_0(1+3x,9)=3$, by the analogue of Lemma 4.2 at depths 4 and 5 together with an explicit member of $T_4(1+2x,4)$. Here too the values are the source's and only the certification is ours; Proposition 10.1 certifies only the halves ≥ 3.

What is *not* settled is Remark 5.3 itself. Section 11 records how far the elementary argument reaches — any further exception has $b_0 \leq a_0^2$ — and exhibits two sub-families on which the inequalities used here are provably too lossy, which is, we suspect, what the authors meant by "rather technical".

Every numbered statement below except those explicitly labelled as computations outside the formal development is machine-checked in Lean 4 [2]; see Section 12.

2. Preliminaries

All polynomials have integer coefficients. We write f_i for the coefficient of x^i in f, so that $\|f\| = \sum_i f_i^2$, and $f|_m = \sum_{i<m} f_i x^i$.

Definition 2.1 ([6, Definition 4.2])**.** A pair (c,d) with $c,d \in \mathbb{Z}[x]$ and $c(0),d(0)\neq 0$ is *weakly robust* if every factorisation $ab=cd$ in $\mathbb{Z}[x]$ satisfies $\|a\|+\|b\| \geq \|c\|+\|d\|-1$. The pair is *robust* if, in addition, for all such (a,b) the relation $a(x)a(x^{-1})+b(x)b(x^{-1}) = c(x)c(x^{-1})+d(x)d(x^{-1})$ forces $(a,b)=\pm(c,d)$ or $\pm(d,c)$, and $\|a\|+\|b\|=\|c\|+\|d\|-1$ forces $a+b\neq 0$.

Remark 2.2 (robust pairs are weakly robust). Every hypothesis below is the weak one, so the only thing this paper needs from the second half of Definition 2.1 is that it strengthens the first. The source states the two notions in one definition and, having defined weak robustness, continues verbatim: "The pair (c,d) is *robust* if for all (a,b) as above, the additional relation $a(x)a(x^{-1})+b(x)b(x^{-1}) = c(x)c(x^{-1})+d(x)d(x^{-1})$ implies that $(a,b)=\pm(c,d)$ or $\pm(d,c)$, and whenever $\|a\|+\|b\|=\|c\|+\|d\|-1$, we have $a+b\neq 0$." We read the word "additional" as cumulative, which is how Definition 2.1 renders it, and the source's own argument requires that reading: the proof of [6, Lemma 4.5] assumes only that (c,d) is robust and then invokes [6, Corollary 4.4], which is stated for weakly robust pairs. For the family treated here nothing depends on the reading in any case, by the next paragraph.

For the family at hand, [6] states without proof, in the sentence preceding its Lemma 5.2, that $(1+kx,l)$ with $|l|>1$ and $k\neq 0$ "is robust if and only if $|k|\geq |l|/p$, where p is the smallest prime divisor of l". For $k\geq 1$ and $l\geq 2$ that criterion is the hypothesis of Lemma 4.1 below: the largest e admitted there is l/p when l is composite, and there is no such e at all when l is prime. So on this family the source's robust pairs are weakly robust whichever way Definition 2.1 is read. We shall need only that sufficiency half, and only for the weak notion, in the divisor form of Lemma 4.1.

Recall that m is *stable* for (c,d) when every $(a,b)\in T_m(c,d)$ satisfies $ab=cd$. Unwinding the disjunction in [6, Corollary 4.4] — "$ab=cd$ or $\|a\|+\|b\|>\|c\|+\|d\|$" — this is exactly the property of m that the corollary asserts for all large m, and $m_0(c,d)\leq M$ means that every $m>M$ is stable. Note that $T_1(c,d)$ consists of the pairs of constants (a_0,b_0) with $a_0b_0=(cd)(0)$ and $a_0^2+b_0^2\leq \|c\|+\|d\|$, so $m_0(c,d)\geq 1$ unless cd is constant.

Two facts from [6] are used as background only and are not re-proved here. The first is Lemma 5.1 of [6], which determines $m_0(1+kx,1+lx)$ for every k and l: it is 2 if $|k-l|\geq 6$, it is 3 if $2\leq |k-l|\leq 5$ and $\max\{|k|,|l|\}\geq 3$, it is 1 if $|k-l|\leq 1$, with the four exceptional values $m_0(1,1+2x)=4$, $m_0(1+x,1-x)=2$, $m_0(1+2x,1-x)=5$, $m_0(1+2x,1-2x)=4$ and the symmetries $m_0(1+lx,1+kx)=m_0(1+kx,1+lx)=m_0(1-kx,1-lx)$. The second is Lemma 5.2, stated in Section 1. Both were re-verified numerically before any of the work below was attempted;

see Remark 2.3, which also records the one degenerate pair at which Lemma 5.1 needs the source's standing hypothesis $c \neq \pm d$.

Remark 2.3 (a computation outside the formal development). An independent implementation in C of the recursion of [6] — build T_1 from the divisors of $(cd)(0)$, build T_{m+1} from T_m by choosing the new coefficient pair subject to the x^m equation and the running weight budget, stop when every pair of T_m has $ab = cd$ — reproduces, exactly and on its first run, all seven values of m_0 printed in [6]: $m_0(1, 1+2x) = 4$, $m_0(1+x, 1-x) = 2$, $m_0(1+2x, 1-x) = 5$, $m_0(1+2x, 1-2x) = 4$, $m_0(1+2x, 4) = 4$, $m_0(1+3x, 9) = 3$ and $m_0(1, 1-7x^2) = 20$ (the last is the entry at $k = 7$ of the first table of [6, Remark 5.4]). It re-verifies all three cases of [6, Lemma 5.1] over the 1326 pairs with $|k|, |l| \leq 25$ and $k \leq l$, and [6, Lemma 5.2] over the 2578 qualifying pairs with $k, l \leq 60$. There is exactly one disagreement in the two sweeps together, and it is not a counterexample: at $(k, l) = (0, 0)$ the lemma's case $|k - l| \leq 1$ reads $m_0 = 1$ while the recursion returns 0; there $c = d = 1$, which [6] excludes throughout by its standing hypothesis $c \neq \pm d$. These are measurements, not theorems, and nothing below depends on them; they are reported because a family of claims about a published quantity should begin by reproducing the published values of that quantity.

3. The stabilisation lemma

Theorem 3.1. *Let $c, d \in \mathbb{Z}[x]$ with $(cd)(0) \neq 0$, let (c, d) be weakly robust, and let $m \geq 1$. If m is stable for (c, d), then so is $m + 1$.*

Proof. Let $(a, b) \in T_{m+1}(c, d)$ and put $a' = a|_m$, $b' = b|_m$, $\alpha = a_m$, $\beta = b_m$, so that $a = a' + \alpha x^m$ and $b = b' + \beta x^m$.

Step 1: $(a', b') \in T_m(c, d)$. The degrees are $< m$ by construction; the first m coefficients of a product depend only on the first m coefficients of the factors, so $a'b' \equiv ab \equiv cd \bmod x^m$; and $\|a'\| + \|b'\| = \|a\| + \|b\| - \alpha^2 - \beta^2 \leq \|c\| + \|d\|$. By stability of m, therefore, $a'b' = cd$.

Step 2: $a_0\beta + \alpha b_0 = 0$. In $(ab)_m = \sum_{i+j=m} a_i b_j$ exactly two terms involve a coefficient of index m, namely $i = 0$ and $i = m$, and these are absent from $(a'b')_m$; every other term is common to both. Since $(ab)_m = (cd)_m = (a'b')_m$, the two extra terms cancel.

Step 3: $\alpha^2 + \beta^2 \leq 1$. Weak robustness applied to the factorisation $a'b' = cd$ of Step 1 gives $\|a'\| + \|b'\| \geq \|c\| + \|d\| - 1$, while $\|a'\| + \|b'\| + \alpha^2 + \beta^2 = \|a\| + \|b\| \leq \|c\| + \|d\|$.

Step 4: $\alpha = \beta = 0$. From Step 1, $a_0 b_0 = (a'b')_0 = (cd)_0 \neq 0$, so $a_0 \neq 0 \neq b_0$. If $\alpha \neq 0$ then $\alpha^2 = 1$ and $\beta = 0$ by Step 3, and Step 2 gives $\alpha b_0 = 0$, which is impossible. Hence $\alpha = 0$, and then Step 2 reads $a_0\beta = 0$, so $\beta = 0$.

Thus $a = a'$ and $b = b'$, and $ab = cd$ by Step 1. □

Corollary 3.2. *Let $c, d \in \mathbb{Z}[x]$ with $(cd)(0) \neq 0$ and (c, d) weakly robust, and let $M \geq 0$. Then $m_0(c, d) \leq M$ if and only if every pair of the single set $T_{M+1}(c, d)$ satisfies $ab = cd$. The fact underlying this is that stability propagates: if m is stable and $1 \leq m \leq m'$, then m' is stable.*

Proof. The forward implication is immediate from the definition, since $M+1 > M$. For the converse, apply Theorem 3.1 repeatedly, starting from $M + 1 \geq 1$. □

Remark 3.3. The passage from [6, Corollary 4.4] quoted in the introduction observes that "forgetting the x^m term gives a natural map $T_{m+1} \to T_m$", which is what makes T_{m+1} cheap to construct from T_m; that is Step 1 of the proof above. What the paper does not record is that stability propagates *along* that map, which is what licenses stopping the loop at the first stable m. We do not doubt that the authors had an argument in mind — the one above is fifteen lines and uses nothing beyond their own definitions — and we state Theorem 3.1 as the missing justification of a criterion of theirs, not as a competing result. Its role here is practical: without it each of the theorems below would be an assertion about infinitely many sets T_m; with it, each is an assertion about T_3.

4. The pairs $(1 + kx,\, l)$

From here on $c = 1 + kx$ and $d = l$ with $k \geq 1$ and $l \geq 2$, so that $cd = l + klx$ and $\|c\| + \|d\| = 1 + k^2 + l^2$.

Lemma 4.1. *Let $k \geq 1$ and $l \geq 2$ be integers and suppose that $e \leq k$ for every factorisation $l = ef$ with $e \geq 2$ and $f \geq 2$. Then $(1 + kx,\, l)$ is weakly robust.*

Proof. A factorisation $ab = l(1+kx)$ has $\deg a + \deg b = 1$, so after exchanging a and b we may write $a = a_0$ and $b = b_0 + b_1 x$ with $a_0 b_0 = l$ and $b_1 = k b_0$. The required inequality $\|a\| + \|b\| \geq k^2 + l^2$ becomes $a_0^2 + b_0^2(1 + k^2) \geq k^2 + a_0^2 b_0^2$, that is, $(b_0^2 - 1)(k^2 - a_0^2 + 1) + 1 \geq 0$. This is automatic when $|b_0| = 1$ or $|a_0| \leq 1$, and otherwise the hypothesis applied to $l = |a_0| \cdot |b_0|$ gives $|a_0| \leq k$, whence the second factor is positive. □

The hypothesis of Lemma 4.1 is the divisor form of the criterion $k \cdot p \geq l$ stated in [6] for the stronger notion, p being the least prime factor of l; it is stated in divisor form because that is the form in which it is verifiable by a finite test on l alone, which the proof of Theorem 7.1 needs. Lemma 4.1 is elementary and is included only for that reason.

Specialising Corollary 3.2 at $M = 2$ turns the assertion $m_0(1 + kx, l) \leq 2$ into a system of Diophantine conditions in six integers.

Lemma 4.2. *Let $k \geq 1$ and $l \geq 2$. A pair (a, b) lies in $T_3(1 + kx,\, l)$ if and only if its coefficients $a_0, a_1, a_2, b_0, b_1, b_2$ satisfy*

$$a_0 b_0 = l, \tag{2}$$

$$a_0 b_1 + a_1 b_0 = kl, \tag{3}$$

$$a_0 b_2 + a_1 b_1 + a_2 b_0 = 0, \tag{4}$$

$$a_0^2 + a_1^2 + a_2^2 + b_0^2 + b_1^2 + b_2^2 \leq 1 + k^2 + l^2; \tag{5}$$

and for such a pair, $ab = cd$ holds if and only if $a_2 = b_2 = 0$. Consequently, if $(1 + kx, l)$ is weakly robust and every solution of (2)–(5) in $\mathbb{Z}^6$ has $a_2 = b_2 = 0$, then $m_0(1 + kx,\, l) \leq 2$.

Proof. The equivalence with (2)–(5) is the definition (1) written out for $m = 3$, using $cd = l + klx$. If $ab = cd$ then $(ab)_3 = (ab)_4 = 0$, and these read $a_1 b_2 + a_2 b_1 = 0$ and $a_2 b_2 = 0$; combined with (4) they force $a_2 = b_2 = 0$, since $a_0 b_0 = l \neq 0$. Conversely $a_2 = b_2 = 0$ makes a and b linear, and (2)–(4) then say precisely $ab = cd$. The last assertion is Corollary 3.2 with $M = 2$, whose hypothesis $(cd)(0) = l \neq 0$ holds. □

Throughout, we call a solution of (2)–(5) with $(a_2, b_2) \neq (0, 0)$ a *bad pair* for (k, l). By Lemma 4.2, Remark 5.3 of [6] asserts that a weakly robust $(1+kx, l)$ has no bad pair unless $(k, l) \in \{(2, 4), (3, 9)\}$.

5. The trivial splitting, and every prime l

Equation (2) splits l as $a_0 b_0$. We call the splitting *trivial* when $|a_0| = 1$ or $|b_0| = 1$. The following theorem disposes of every trivial splitting at once, for all k and all l, and with no computation.

Theorem 5.1. *Let $k \geq 1$ and $l \geq 2$ be integers. Then no bad pair for (k, l) has $|a_0| = 1$ or $|b_0| = 1$.*

Proof. The system (2)–(5) is invariant under exchanging (a_0, a_1, a_2) with (b_0, b_1, b_2) and under negating all six variables, so it suffices to treat $a_0 = 1$, and then $b_0 = l$. Equation (3) reads $b_1 + l a_1 = kl$, so $l \mid b_1$; write $b_1 = l\beta$, whence $a_1 = k - \beta$. Equation (4) reads $b_2 + a_1 b_1 + a_2 l = 0$, so $b_2 = -l\gamma$ with $\gamma = (k - \beta)\beta + a_2$, i.e. $a_2 = \gamma - (k - \beta)\beta$. Substituting into (5) and cancelling $1 + l^2$ gives

$$(k - \beta)^2 + l^2\beta^2 + a_2^2 + l^2\gamma^2 \leq k^2, \qquad \text{that is} \qquad \beta^2(1 + l^2) + a_2^2 + l^2\gamma^2 \leq 2k\beta.$$

If $\beta < 0$ the right-hand side is negative and the left-hand side is not. If $\beta = 0$ the inequality reads $a_2^2 + l^2\gamma^2 \leq 0$, forcing $a_2 = \gamma = 0$ and hence $a_2 = b_2 = 0$, so the pair is not bad. So assume $\beta \geq 1$.

Using $2k\beta = 2(k-\beta)\beta + 2\beta^2 = 2(\gamma - a_2) + 2\beta^2$ together with $l \geq 2$, i.e. $l^2 - 1 \geq 3$ and $l^2 \geq 4$, we obtain $3\beta^2 + a_2^2 + 4\gamma^2 \leq 2\gamma - 2a_2$; multiplying by 4 and completing squares,

$$12\beta^2 + 4(a_2+1)^2 + (4\gamma - 1)^2 \leq 5,$$

which contradicts $\beta \geq 1$. □

Theorem 5.2. *For every integer $k \geq 1$ and every prime l,*

$$m_0(1+kx,\, l) \leq 2.$$

Proof. A prime admits no factorisation $l = ef$ with $e, f \geq 2$, so the hypothesis of Lemma 4.1 holds vacuously and $(1+kx, l)$ is weakly robust. For the same reason, (2) forces $|a_0| = 1$ or $|b_0| = 1$, so by Theorem 5.1 there is no bad pair, and Lemma 4.2 gives the conclusion. □

Theorem 5.2 is an infinite sub-family of Remark 5.3, proved rather than tested, and no part of its proof is a computation. It is worth comparing its range with what [6] establishes. For prime l the least prime factor is l itself, so the source's robustness criterion $|k| \geq |l|/p$ holds for every $k \geq 1$: all the pairs of Theorem 5.2 are robust, and those of them that Lemma 5.2 does not cover are pairs Remark 5.3 speaks of. Lemma 5.2 of [6] proves the stronger conclusion $m_0 = 1$, but only when $l \geq 5$ *and* $k \leq l^2/2$. So $l = 2$, $l = 3$, and every pair with $k > l^2/2$, lie outside what [6] proves and inside Theorem 5.2.

Remark 5.3 (a computation outside the formal development). A direct enumeration of $T_3(1+kx, l)$ over every pair with $k, l \leq 400$ passing the criterion $kp \geq l$ of [6] recalled in Section 2 (p the least prime factor of l; such pairs are weakly robust by Lemma 4.1) finds exactly two cells carrying a bad pair, $(k,l) = (2,4)$ and $(k,l) = (3,9)$, with witnesses splitting l as $a_0 = b_0 = 2$ and as $a_0 = b_0 = 3$. The enumeration visits the splittings in increasing a_0, so it also shows that no bad pair in that range splits l trivially, which is the observation Theorem 5.1 turns into a theorem for all k and l.

6. Narrowing the non-trivial splitting

What remains after Theorem 5.1 is $l = a_0 b_0$ with $2 \leq a_0 \leq b_0$ (signs and order are normalised by the symmetries used in its proof). Two inequalities, both uniform in k and l, cut that region down to a two-parameter band.

Lemma 6.1. *Let $a_0, a_1, a_2, b_0, b_1, b_2$ be a bad pair for (k,l) with $2 \leq a_0 \leq b_0$. Then*

$$k^2 \leq a_0^2 + b_0^2.$$

Proof. Put $s = a_0^2 + b_0^2$, $u = a_1^2 + b_1^2$, $v = a_2^2 + b_2^2$. Lagrange's identity and (3) give

$$s\,u = (a_0 b_1 + a_1 b_0)^2 + (a_0 a_1 - b_0 b_1)^2 \geq (kl)^2,$$

and $v \geq 1$ because the pair is bad. Multiplying (5) by $s > 0$ and substituting both bounds yields $s^2 + k^2 l^2 + s \leq s(1 + k^2 + l^2)$, that is, $(s - k^2)(s - l^2) \leq 0$. Finally $s = a_0^2 + b_0^2 < a_0^2 b_0^2 = l^2$ because $a_0, b_0 \geq 2$, so the second factor is negative and $s \geq k^2$. □

Lemma 6.2. *Let $k \geq 1$ and let $(1+kx,\, l)$ be weakly robust. Then every factorisation $l = a_0 b_0$ with $2 \leq a_0 \leq b_0$ satisfies*

$$b_0 \leq k.$$

Proof. Apply Definition 2.1 to the factorisation $l(1+kx) = b_0 \cdot \big(a_0(1+kx)\big)$: it gives $b_0^2 + a_0^2 + a_0^2 k^2 \geq k^2 + l^2$, that is, $k^2(a_0^2 - 1) \geq a_0^2 b_0^2 - a_0^2 - b_0^2$. Suppose $k \leq b_0 - 1$. Then $k^2 \leq (b_0 - 1)^2$, and substituting gives $(a_0^2 - 1)(2b_0 - 1) \leq a_0^2$. But $(a_0^2 - 1)(2b_0 - 1) \geq (a_0^2 - 1)(2a_0 - 1) = 2a_0^3 - a_0^2 - 2a_0 + 1 > a_0^2$ for $a_0 \geq 2$, a contradiction. □

Together, a bad pair for a weakly robust $(1 + kx, l)$ with non-trivial splitting satisfies

$$2 \le a_0 \le b_0 \le k \le \sqrt{a_0^2 + b_0^2} \le b_0\sqrt{2}, \qquad l = a_0 b_0. \tag{6}$$

Neither inequality mentions a_1, a_2, b_1, b_2: they are conditions on the divisor pair alone. That is what makes them useful, since it means the outer three loops of any search — over k, a_0, b_0 — can be pruned before the inner enumeration starts.

Remark 6.3 (a measurement, outside the formal development). Over the 52 118 pairs with $k, l \le 250$ passing the same criterion $kp \ge l$, the triples (k, a_0, b_0) surviving (6) number 235, and the resulting enumeration over (a_1, a_2) costs 2.5×10^5 inner steps. The same enumerator with the two inequalities of (6) deleted costs 3.6×10^7, and building T_1, T_2, T_3 outright by the recursion of Remark 2.3 costs 2.1×10^8. This is a statement about three programs, not about the mathematics, and nothing below depends on it; it is recorded because it is the reason the computation of Section 7 is affordable at all.

Remark 6.4 (quantifying a narrowing lemma). Lemma 6.2 concludes $b_0 \le k$ for *every* factorisation of l into two factors ≥ 2, not merely for the one appearing in a given bad pair, and the difference is not cosmetic. Consider $(k, l) = (4, 16)$ with

$$a = 4 + 2x - 4x^2, \qquad b = 4 + 14x - 3x^2,$$

a bad pair: (2)–(4) read $16 = 16$, $4 \cdot 14 + 2 \cdot 4 = 64 = kl$ and $4 \cdot (-3) + 2 \cdot 14 + (-4) \cdot 4 = 0$, and the weight is $257 \le 273 = 1 + k^2 + l^2$. Its own splitting is $a_0 = b_0 = 4$, which satisfies both (6) inequalities ($b_0 = 4 \le k = 4$ and $k^2 = 16 \le 32$). Likewise $(k, l) = (9, 81)$ carries the bad pair $a = 9 + 3x - 13x^2$, $b = 9 + 78x - 13x^2$ of weight $6593 \le 6643$, with splitting $a_0 = b_0 = 9$. Neither pair is weakly robust: the factorisations $16 = 2 \cdot 8$ and $81 = 3 \cdot 27$ violate Lemma 6.2, since $8 > 4$ and $27 > 9$. So both are excluded by the hypothesis of Theorem 7.1, but only if the lemma is used at all factorisations of l; used at the bad pair's own splitting it excludes neither. We record this because a search organised around the weaker reading returns the wrong answer at exactly these two cells, and for no others in the range we examined.

7. Remark 5.3 over a box

Theorem 7.1. *Let k and l be integers with $1 \le k \le 1000$ and $2 \le l \le 1000$, let $(1 + kx,\, l)$ be weakly robust, and suppose $(k, l) \ne (2, 4)$ and $(k, l) \ne (3, 9)$. Then*

$$m_0(1 + kx,\, l) \le 2.$$

Proof sketch. By Lemma 4.2 it suffices to show that (k, l) has no bad pair. Suppose one exists. By Theorem 5.1 its splitting is non-trivial, and after the symmetries of that theorem we may assume $2 \le a_0 \le b_0$ with $a_0 b_0 = l$. Lemmas 6.1 and 6.2 then place (k, a_0, b_0) in the band (6). Moreover Lemma 6.2, applied at *every* factorisation of l and not only at the one carried by the bad pair, forces $\max\{e, l/e\} \le k$ whenever $l = e \cdot (l/e)$ with both factors ≥ 2; this is a finite test on k and l alone, and Remark 6.4 shows it cannot be dispensed with.

What remains is a finite claim, and it is settled by exhaustive computation. For each triple (a_0, b_0, k) of integers with $2 \le a_0 \le b_0 \le k \le 1000$, $a_0 b_0 \le 1000$, $k^2 \le a_0^2 + b_0^2$, $(k, a_0 b_0) \notin \{(2, 4), (3, 9)\}$, and passing the divisor test just described, the following enumeration is run. Put $l = a_0 b_0$ and $R_0 = 1 + k^2 + l^2 - a_0^2 - b_0^2$. For each integer a_1 with $a_1^2 \le R_0$, equation (3) determines $b_1 = (kl - a_1 b_0)/a_0$ when a_0 divides the numerator and rules the value out otherwise; put $R_1 = R_0 - a_1^2 - b_1^2$ and discard the value if $R_1 < 0$. For each integer a_2 with $a_2^2 \le R_1$, equation (4) determines $b_2 = (-a_1 b_1 - a_2 b_0)/a_0$ in the same way, and the triple fails the test if $a_2^2 + b_2^2 \le R_1$ and $(a_2, b_2) \ne (0, 0)$. Every triple passes. Since (5) bounds a_1^2 by R_0 and a_2^2 by R_1, and (3) and (4) determine b_1 and b_2 from the remaining data, the enumeration visits every candidate bad pair with the given constant terms; so no bad pair exists, and the supposition is refuted. □

Three features of the statement are worth isolating. First, the hypothesis is weak robustness, which is implied by the robustness of Remark 5.3 (Remark 2.2), so Theorem 7.1 covers every pair the remark speaks of in the given range and more. Second, it is stated for *all* weakly robust pairs of the box, not only for those that [6, Lemma 5.2] misses; on the pairs that lemma does cover it re-derives the weaker conclusion $m_0 \le 2$ in place of $m_0 = 1$. Third, the two exception clauses are necessary, by Proposition 10.1 below. As a measurement rather than a mathematical claim: the enumeration described in the proof takes about 29 seconds of processor time, and the module containing it about 67, with a memory footprint that is the cost of loading the library and not of the search.

Corollary 7.2. $m_0(1 + 3x,\, 6) \le 2$.

Proof. The factorisations of 6 into two factors ≥ 2 are $2 \cdot 3$ and $3 \cdot 2$, and in each the first factor is at most $3 = k$, so $(1 + 3x, 6)$ is weakly robust by Lemma 4.1; the pair is not an exception, and $k, l \le 1000$. Apply Theorem 7.1. □

Corollary 7.2 is a single instance of Theorem 7.1 and is recorded only to show the hypotheses are satisfiable outside the source's proved range: the least prime factor of 6 is $p = 2$, so $kp = 6 \ge l$ and the pair is robust, while [6, Lemma 5.2] requires $k > \sqrt{p^2 + l^2/p^2} = \sqrt{13}$, which $k = 3$ fails.

8. Each level of the search is a line meeting a disk

What keeps Theorem 7.1 inside a box is the shape of its two inner loops, and that is a question with a clean answer. Once Theorem 5.1 and (6) have fixed a triple (a_0, b_0, k) with $l = a_0 b_0$, what is left of the system (2)–(5) is two levels of the same shape — one linear equation in two unknowns, under one quadratic budget:

$$\text{level 1:}\quad a_0 b_1 + b_0 a_1 = kl,\quad a_1^2 + b_1^2 \le R_0; \qquad \text{level 2:}\quad a_0 b_2 + b_0 a_2 = -a_1 b_1,\quad a_2^2 + b_2^2 \le R_1,$$

where $R_0 = 1 + k^2 + l^2 - a_0^2 - b_0^2$ and $R_1 = R_0 - a_1^2 - b_1^2$, by (3), (4) and (5). Each level asks for the lattice points where a line meets a disk, and the following inequality — one Lagrange identity, the same one as in Lemma 6.1 — says how long the chord is.

Lemma 8.1 (the chord bound)**.** *Let a_0, b_0 be positive integers, put $s = a_0^2 + b_0^2$, and let $u, R \in \mathbb{Z}$. If $x, y \in \mathbb{Z}$ satisfy*

$$a_0 y + b_0 x = u \qquad \text{and} \qquad x^2 + y^2 \le R,$$

then

$$(a_0 x - b_0 y)^2 \le sR - u^2 \qquad \text{and} \qquad \bigl|sx - b_0 u\bigr| \le a_0 \sqrt{sR - u^2}. \tag{7}$$

In particular there is no such pair when $sR < u^2$, and otherwise every such x lies within $a_0\sqrt{sR - u^2}/s$ of $b_0 u/s$.

Proof. Lagrange's identity reads $(a_0 x - b_0 y)^2 + (a_0 y + b_0 x)^2 = s(x^2 + y^2)$; substituting the equation on the left and the budget on the right gives the first inequality, and with it $sR - u^2 \ge 0$. Substituting the equation into $sx - b_0 u$ gives $sx - b_0 u = (a_0^2 + b_0^2)x - a_0 b_0 y - b_0^2 x = a_0(a_0 x - b_0 y)$, and the second inequality follows from the first. □

Both levels are instances: level 1 with $(x, y) = (a_1, b_1)$, $u = kl$, $R = R_0$, and level 2 with $(x, y) = (a_2, b_2)$, $u = -a_1 b_1$, $R = R_1$. So each level is enumerated over an interval of about $2a_0\sqrt{sR - u^2}/s$ integers centred at $b_0 u/s$ — in the implementation, over the integers within $\lfloor a_0 \lfloor \sqrt{sR - u^2} \rfloor / s \rfloor + 1$ of $\lfloor b_0 u/s \rfloor$, the extra unit absorbing both roundings — and is skipped altogether when $sR < u^2$.

Three features of these windows are worth naming. First, the level-1 discriminant is

$$sR_0 - (kl)^2 = s + (l^2 - s)(s - k^2), \tag{8}$$

which is exactly the right-hand side of the displayed inequality of Remark 11.1: the quantity that the hand argument of Section 11 produces as a budget is here the width of a loop. By Lemma 6.1 it

is non-negative throughout the band (6), so level 1 is never empty for that reason; and since $l^2 > s$ it decreases as k rises, so it is smallest at the top of the k-interval, where it collapses to s itself — and the window to at most two integers — exactly when s is a square and $k = \sqrt{s}$. Second, $k \geq b_0$ gives $s - k^2 \leq a_0^2$, so (8) is at most $s + l^2 a_0^2$ and the level-1 window holds $O(a_0^3/b_0)$ integers, where the budget window $|a_1| \leq \sqrt{R_0}$ used in the proof of Theorem 7.1 holds $O(l)$. Third, at level 2 the gain comes mostly from the emptiness test rather than from the width: $u = -a_1 b_1$ is usually far larger in absolute value than $\sqrt{sR_1}$, so the level is skipped without any enumeration at all (Remark 9.5).

We claim nothing new for Lemma 8.1. It is recorded because it is the reason the range of the next section is affordable, and because a window derived from the level's own equation is sound by construction: unlike a hand-tuned window it cannot silently omit a candidate.

9. Remark 5.3 for every k, and for $2 \leq l \leq 4 \cdot 10^4$

Theorem 9.1. *Let k and l be integers with $k \geq 1$ and $2 \leq l \leq 4 \cdot 10^4$, let $(1 + kx,\, l)$ be weakly robust, and suppose $(k, l) \neq (2, 4)$ and $(k, l) \neq (3, 9)$. Then*

$$m_0(1 + kx,\, l) \leq 2.$$

Proof sketch. By Lemma 4.2 it suffices to show that (k, l) has no bad pair. Suppose one exists. By Theorem 5.1 its splitting is non-trivial, and after the symmetries of that theorem we may assume $l = a_0 b_0$ with $2 \leq a_0 \leq b_0$; then $a_0^2 \leq l$, and Lemmas 6.1 and 6.2 place (k, a_0, b_0) in the band (6). So with $L = 4 \cdot 10^4$ the triples to be examined are exactly

$$2 \leq a_0 \leq \lfloor \sqrt{L} \rfloor, \qquad a_0 \leq b_0 \leq \lfloor L/a_0 \rfloor, \qquad b_0 \leq k \leq \Big\lfloor \sqrt{a_0^2 + b_0^2} \Big\rfloor, \tag{9}$$

of which there are 955 002. The hypothesis $k \leq 1000$ of Theorem 7.1 has no counterpart here because k is enumerated over the interval that the band leaves — of length about $a_0^2/2b_0$ — rather than over a side of a box.

What remains is a finite claim, and it is settled by exhaustive computation. Each triple (9) is disposed of in one of three ways. It may be one of the two exceptions. Or l may have a divisor e with $2 \leq e < a_0$ and $ek < l$; then $l = e \cdot (l/e)$ has both factors ≥ 2 and $e < a_0 \leq b_0 < l/e$, so if $(1 + kx, l)$ were weakly robust, Lemma 6.2 applied to that factorisation would give $l/e \leq k$, that is $l \leq ek$, against $ek < l$. The triple is therefore excluded by the hypothesis. (Divisors $e \geq a_0$ with cofactor $\geq e$ need not be tested: there $l/e \leq l/a_0 = b_0 \leq k$ already, so the criterion cannot fail. This shortens the divisor loop from $O(b_0)$ steps to $O(a_0)$.) Otherwise the two levels of Section 8 are enumerated. For each a_1 in the level-1 window, equation (3) determines $b_1 = (kl - a_1 b_0)/a_0$ when a_0 divides the numerator and rules the value out otherwise; then $R_1 = R_0 - a_1^2 - b_1^2$ is formed and the value discarded if $R_1 < 0$. For each a_2 in the level-2 window, (4) determines $b_2 = (-a_1 b_1 - a_2 b_0)/a_0$ in the same way, and the triple fails the test if $a_2^2 + b_2^2 \leq R_1$ and $(a_2, b_2) \neq (0, 0)$. At $L = 4 \cdot 10^4$ exactly 64 926 of the triples are neither one of the two exceptions nor excluded by the divisor test; each of those is handed to the enumeration, and each of them passes. By Lemma 8.1 each window contains every integer solution of its level, so the enumeration visits every candidate bad pair with the given constant terms; hence no bad pair exists, and the supposition is refuted. □

Theorem 9.1 contains Theorem 7.1, since $1000 \leq 4 \cdot 10^4$ and no condition on k is imposed; the formal development records that derivation, so the older statement has a second proof that does not use the older computation. The containment is strict in both directions, and the following two instances say by how much.

Proposition 9.2. *$m_0(1 + 20000\,x,\, 39999) \leq 2$ and $m_0(1 + 10^6 x,\, 6) \leq 2$.*

Proof. Both are instances of Theorem 9.1, and neither pair is an exception. For $l = 39999 = 3 \cdot 67 \cdot 199$ the factorisations $l = ef$ with $e, f \geq 2$ have $e \in \{3, 67, 199, 201, 597, 13333\}$, all at most $k = 20000$, so $(1 + 20000x, 39999)$ is weakly robust by Lemma 4.1; and $l \leq 4 \cdot 10^4$. For $l = 6$ the factorisations are $2 \cdot 3$ and $3 \cdot 2$, and $2, 3 \leq 10^6 = k$. □

The first instance sits at forty times the l-reach of Theorem 7.1, but it is *not* new relative to [6]: all three hypotheses of [6, Lemma 5.2] hold at that pair. Indeed $l = 39999 \geq 5$; its least prime factor is $p = 3$, so that $\sqrt{p^2 + l^2/p^2} = \sqrt{9 + 13333^2} < 13334 < 20000 = k$; and $k = 20000 \leq l^2/2 = 799\,960\,000.5$. So that lemma covers the pair, and with the stronger conclusion $m_0(1 + 20000\,x,\, 39999) = 1$. What Proposition 9.2 contributes there is reach — an instance of Theorem 9.1 far outside the box of Theorem 7.1 — and not new mathematics. The second instance has k a thousand times the k-reach of Theorem 7.1, and Lemma 5.2 does *not* cover it: that lemma needs $k \leq l^2/2 = 18$. It is nonetheless a pair the source calls robust, since $kp = 2 \cdot 10^6 \geq 6 = l$.

Proposition 9.3. (i) *The two-level test described in the proof of Theorem 9.1 returns* false *at each of the three triples* $(k, a_0, b_0) = (2, 2, 2)$, $(3, 3, 3)$ *and* $(4, 4, 4)$.
(ii) $(1 + 4x,\, 16)$ *is not weakly robust, and* $m_0(1 + 4x,\, 16) \geq 3$.

Proof. (i) Three evaluations of the test. Each of the three triples carries a bad pair inside its windows: the two witnesses of Proposition 10.1 at $(2, 2, 2)$ and $(3, 3, 3)$, and the pair displayed in Remark 6.4 at $(4, 4, 4)$. (ii) The factorisation $16 = 2 \cdot 8$ has both factors ≥ 2 and $8 > 4 = k$, so Lemma 6.2 denies weak robustness. The same displayed pair, $a = 4 + 2x - 4x^2$, $b = 4 + 14x - 3x^2$, lies in $T_3(1 + 4x, 16)$ and has $(ab)_3 \neq 0$, so 3 is not stable and $m_0(1 + 4x, 16) \geq 3$. □

Part (i) is what stops Theorem 9.1 from being vacuously true. A search whose windows were empty — the failure mode a narrowing lemma invites — would return a passing verdict everywhere, and in particular at these three triples, where a passing verdict is false. Part (ii) is the corresponding statement about the hypotheses: $(k, l) = (4, 16)$ lies well inside the range of Theorem 9.1 and is not one of the two exceptions, yet its conclusion fails there. So the weak-robustness hypothesis cannot be dropped, and the restriction of Remark 5.3 to robust pairs is necessary rather than decorative.

Remark 9.4 (a computation outside the formal development). The windows of Section 8 were checked against brute force before being used. An independent program in C enumerates the solution set of (3)–(5) at every triple (9) twice, once with the budget windows $|a_1| \leq \sqrt{R_0}$, $|a_2| \leq \sqrt{R_1}$ of the proof of Theorem 7.1 and once with the chord windows, and compares an order-independent hash of the two multisets of solutions; they are identical at $L = 50$, 120, 250, 600 and 1200. Run over the whole of (9) with $L = 4 \cdot 10^4$, and without the two skips that the proof of Theorem 9.1 makes, the chord enumeration finds bad pairs at exactly four cells: $(k, l) = (2, 4)$ with 3 solutions, $(3, 9)$ with 10, $(4, 16)$ with 22 and $(9, 81)$ with 22. The first two are the source's exceptions and the last two are not weakly robust — $16 = 2 \cdot 8$ and $81 = 3 \cdot 27$ — which is what Theorem 9.1 predicts and Proposition 9.3 illustrates at three of the four. A window that dropped solutions would show up here as a smaller count.

Remark 9.5 (a measurement, outside the formal development). The evaluation in the proof of Theorem 9.1 was run at a range of bounds L, with these step counts (iterations of the level-1 loop, from the C mirror) and processor times (Lean's compiled evaluator, with the cost of loading the library subtracted):

L	10^3	$2 \cdot 10^3$	$4 \cdot 10^3$	10^4	$2 \cdot 10^4$	$4 \cdot 10^4$
steps	$6.9 \cdot 10^4$	$3.1 \cdot 10^5$	$1.4 \cdot 10^6$	$1.09 \cdot 10^7$	$5.11 \cdot 10^7$	$2.59 \cdot 10^8$
seconds	≈ 0	0.5	2.1	18.1	83.8	405.2

The step count grows like $L^{2.25}$ and the cost per step is a flat 1.6 microseconds, so the extrapolation over this decade is good to a few percent. The module containing the $L = 4 \cdot 10^4$ evaluation checks in 444 seconds and the module of Proposition 10.3 in 21; the memory footprint of both is the cost of loading the library. Two independent obstacles stop the same method at $L = 10^5$: about $2 \cdot 10^9$ steps, some 52 minutes for a single evaluation; and the level-1 discriminant (8), which grows like $L^4/4$ and leaves the evaluator's machine-integer range (bounded by 2^{62}) near $L \approx 6.6 \cdot 10^4$, past which every discriminant is an arbitrary-precision allocation and the figure of 1.6 microseconds is itself an underestimate. Two further measurements, for the record: with the budget windows of the

proof of Theorem 7.1 in place of the chord windows the same enumeration costs $1.77\cdot10^7$ steps at $L=10^3$ and $1.12\cdot10^8$ at $L=2\cdot10^3$, of which $1.66\cdot10^7$ and $1.08\cdot10^8$ — that is 94% and 96% — are spent at the second level; and at $L=10^4$, of the $1.54\cdot10^6$ level-1 candidates that pass the divisibility test, $1.52\cdot10^6$ are discarded by the emptiness test $sR_1<u^2$ alone and only 140 ever enter a level-2 walk, for 7 496 level-2 steps in all. These are statements about programs, not about the mathematics, and nothing above depends on them.

10. The two exceptions, and the constant

The following two propositions certify lower bounds that [6] states as exact values. The mathematics is the source's; what is ours is only the certification, and only of the inequality $\geq$.

Proposition 10.1. *$m_0(1+2x,\,4)\geq3$ and $m_0(1+3x,\,9)\geq3$.*

Proof. For $(k,l)=(2,4)$ take $a=2+2x-2x^2$ and $b=2+2x$. Then $a_0b_0=4=l$, $a_0b_1+a_1b_0=2\cdot2+2\cdot2=8=kl$, $a_0b_2+a_1b_1+a_2b_0=0+4-4=0$, and the weight is $12+8=20\leq21=1+k^2+l^2$, so $(a,b)\in T_3$; but $ab=4+8x-4x^3\neq4+8x=cd$. Hence 3 is not stable, and $m_0\geq3$.

For $(k,l)=(3,9)$ take $a=3+3x-5x^2$ and $b=3+6x-x^2$. Then $a_0b_0=9=l$, $a_0b_1+a_1b_0=18+9=27=kl$, $a_0b_2+a_1b_1+a_2b_0=-3+18-15=0$, and the weight is $43+46=89\leq91$, so $(a,b)\in T_3$; but $(ab)_3=a_1b_2+a_2b_1=-3-30=-33\neq0=(cd)_3$. □

Proposition 10.2. *$m_0(1+3x,\,2)\geq2$; in particular the constant 2 in Theorems 5.2 and 7.1 cannot be replaced by 1.*

Proof. Take $a=1+2x$ and $b=2+2x$. Then $ab\equiv2+6x \bmod x^2$, which is cd, the degrees are <2, and the weight is $5+8=13\leq14=1+9+4$; so $(a,b)\in T_2(1+3x,2)$. But $ab=2+6x+4x^2\neq2+6x=cd$, so 2 is not stable. The pair $(1+3x,2)$ is weakly robust by Lemma 4.1 (there is no factorisation of 2 into two factors ≥2), it is covered by Theorem 5.2 since 2 is prime, and it lies in the box of Theorem 7.1. □

The exact values $m_0(1+2x,4)=4$ and $m_0(1+3x,9)=3$ are [6, Remark 5.3] as well. Proposition 10.1 certifies only the halves ≥3; the values themselves require the analogue of Lemma 4.2 at depths 4 and 5, in eight and ten integers, and that is the next proposition. Once more the mathematics is the source's and what is ours is the certification.

Proposition 10.3. *$m_0(1+2x,\,4)=4$ and $m_0(1+3x,\,9)=3$.*

Proof sketch. Both pairs are weakly robust by Lemma 4.1: the only factorisation of 4 into two factors ≥2 is $2\cdot2$ with $2\leq k=2$, and that of 9 is $3\cdot3$ with $3\leq k=3$. So Corollary 3.2 applies, and turns $m_0(1+3x,9)\leq3$ into the single condition that every pair of $T_4(1+3x,9)$ satisfies $ab=cd$, and $m_0(1+2x,4)\leq4$ into the same condition on $T_5(1+2x,4)$. Exactly as in Lemma 4.2, membership in $T_m(1+kx,l)$ is the system of the m bilinear equations $\sum_{i+j=n}a_ib_j=(cd)_n$ for $n<m$ together with the weight inequality $\sum_i a_i^2+\sum_i b_i^2\leq1+k^2+l^2$ in the $2m$ coefficients; and such a pair has $ab=cd$ as soon as every coefficient of index ≥2 of a and of b vanishes, by the argument given for $m=3$ — the vanishing makes a and b linear, and the index-2 congruence then forces $a_1b_1=0$, which is exactly what makes their product $l+klx$. Each of the two systems is then settled by exhaustive computation: the weight inequality confines every coordinate to $x^2\leq1+k^2+l^2$, that is $|x|\leq9$ for $(3,9)$ and $|x|\leq4$ for $(2,4)$, and within those ranges each level's equation determines the corresponding b-coefficient by division, as in the proof of Theorem 9.1. Both enumerations report no solution with a non-vanishing coefficient of index ≥2.

For the lower bounds, $m_0(1+3x,9)\geq3$ is Proposition 10.1, and $m_0(1+2x,4)\geq4$ needs a member of $T_4(1+2x,4)$ with $ab\neq cd$. Take $a=b=2+2x-x^2+x^3$: the four equations read $4=4$, $2\cdot2+2\cdot2=8=kl$, $2(-1)+2\cdot2+(-1)2=0$ and $2\cdot1+2(-1)+(-1)2+1\cdot2=0$, the weight is $2(4+4+1+1)=20\leq21=1+k^2+l^2$, and $(ab)_6=1\cdot1=1\neq0=(cd)_6$. So 4 is not stable for $(1+2x,4)$. □

Similarly $m_0(1 \pm 3x, \pm 2) = 2$ is stated in [6, Example 7.5] — the trinomial example, where the source writes "This also holds for $2 \le |l| \le 4$ with the exceptions $m_0(1 \pm 3x, \pm 2) = 2$ and $m_0(1 \pm 3x, \pm 4) = 2$" — which is where the constant of Proposition 10.2 is seen to be exactly right.

11. What remains

Remark 5.3 of [6] is not proved here; what is proved is an infinite sub-family of it (Theorem 5.2) and a range (Theorems 7.1 and 9.1), together with the two values it names (Proposition 10.3). This section records how far the elementary argument reaches and why it stops, in the hope of saving the next attempt some time.

Remark 11.1 (proved on paper; not part of the formal development). Keep $2 \le a_0 \le b_0$, $l = a_0 b_0$, $s = a_0^2 + b_0^2$, and let a bad pair be given for a weakly robust $(1 + kx, l)$. Write $g = \gcd(a_0, b_0)$, $a_0 = g\alpha$, $b_0 = g\beta$ with $\gcd(\alpha, \beta) = 1$. Then (3) reads $\alpha b_1 + \beta a_1 = kg\alpha\beta$, so $\alpha \mid a_1$ and $\beta \mid b_1$; write $a_1 = \alpha X$, $b_1 = \beta Y$, and then $X + Y = kg$. Hence $a_1 b_1 = \alpha\beta XY$ and, since $X + Y = kg > 0$, either $XY = 0$ or $|XY| \ge kg - 1$. The case $XY = 0$ is closed by Lemma 6.2, which gives $b_0 \le k$: if $X = 0$ then $a_1 = 0$ and $b_1 = kb_0$, and (5) together with $a_2^2 + b_2^2 \ge 1$ gives $k^2(b_0^2 - 1) \le a_0^2(b_0^2 - 1) - b_0^2$, whence $k^2 < a_0^2 \le b_0^2 \le k^2$; if $Y = 0$ then $b_1 = 0$ and $a_1 = ka_0$, and the same computation gives $k^2(a_0^2 - 1) \le b_0^2(a_0^2 - 1) - a_0^2$, whence $k^2 < b_0^2 \le k^2$. Otherwise, since $g \mapsto l(kg - 1)/g^2$ is decreasing for $kg > 2$ and $g \le a_0$,

$$|a_1 b_1| = \frac{l}{g^2}\,|XY| \;\ge\; \frac{l(kg-1)}{g^2} \;\ge\; b_0\Big(k - \tfrac{1}{2}\Big).$$

Multiplying (5) by s and applying Lagrange's identity twice — once to (3) as in Lemma 6.1, and once to (4) in the form $(a_0 b_2 + a_2 b_0)^2 + (a_0 a_2 - b_0 b_2)^2 = s(a_2^2 + b_2^2)$ with $a_0 b_2 + a_2 b_0 = -a_1 b_1$ — gives the exact inequality

$$(a_0 a_1 - b_0 b_1)^2 + (a_1 b_1)^2 \;\le\; s + (s - k^2)(l^2 - s).$$

By Lemma 6.2, $k \ge b_0$, so $s - k^2 \le a_0^2$; combining with the lower bound on $|a_1 b_1|$, dividing by b_0^2 and using $a_0 \le b_0$ yields $(b_0 - \tfrac{1}{2})^2 \le a_0^4 - a_0^2 + 2$, hence

$$b_0 \le a_0^2.$$

Remark 11.1 does not finish Remark 5.3, and the rest of this paragraph, like the remark itself, is analysis carried out on paper and not part of the formal development. The surviving region $2 \le a_0 \le b_0 \le a_0^2$, $b_0 \le k \le \sqrt{a_0^2 + b_0^2}$, intersected with the robustness condition $kp \ge l$, is still infinite, and it contains two sub-families that no inequality above excludes. The first is the diagonal $a_0 = b_0 = t$: robustness forces $t^2/p \le k \le t\sqrt{2}$ with $p \mid t$, hence $t \le p\sqrt{2}$, hence $t = p$ is prime, so that $l = p^2$ and $p \le k \le \lfloor p\sqrt{2} \rfloor$ — a family containing the two known exceptions $(p, k) = (2, 2)$ and $(3, 3)$. The second is the strip $a_0 = p$, $b_0 = k = l/p$. Every member of either with $l \le 4 \cdot 10^4$ is weakly robust, by Lemma 4.1 and the condition $kp \ge l$, hence settled by Theorem 9.1: on the diagonal that is every prime $p \le 199$, the two exceptions apart. Both families are infinite beyond that. Moreover, on the diagonal the Cauchy–Schwarz relaxation used in Lemma 6.1 is provably too lossy: with $a_0 = b_0 = k = t$ the left-hand side of the displayed inequality of Remark 11.1 has real minimum $t^6 - 4t^4$, strictly below the available budget $t^6 - 2t^4 + 2t^2$. So no argument using only (3), (4), (5) and Cauchy–Schwarz can close the diagonal; what closes it, in the cases we checked by hand, is integrality — the real minimiser is integral only when $t^4 - 8t^2$ is a perfect square, hence only when $t^2 - 8$ is, and $t^2 - u^2 = 8$ forces $t = 3$; while the nearest integral configuration $(X, Y) = (t^2 - 2, 2)$ requires $t \mid XY$, i.e. $t \mid 4$. We take this to be the "rather technical" that the authors anticipated.

Remark 11.2 (a measurement, outside the formal development). The bound of Remark 5.3 was checked numerically over all 13 357 906 pairs with $k, l \le 4000$ passing the criterion $kp \ge l$: the only ones with $m_0 > 2$ are $(2, 4)$ and $(3, 9)$. That is a search, not a proof, and it is the kind of statement Theorem 7.1 is meant to replace on its own range; we report it only to say what the numerical evidence beyond the proved box looks like.

The range of that remark is now inside Theorem 9.1. We keep the remark because its enumeration is the recursion of [6] itself, run without any of the narrowing lemmas and without the chord windows, so it is a check on both rather than a consequence of either.

Enlarging the range past Theorem 7.1 turned out, as expected, to be a matter of engineering rather than mathematics. Its enumerator allocates a fresh list of about $2\sqrt{R_1}$ integers for each value of a_1, which dominates the cost once $l \approx 2000$; replacing both walks by index walks and both budget windows by the chord windows of Lemma 8.1 — the level-1 window being the guess $a_1 = ka_0 + t$, $|t| \le a_0\sqrt{R_0}/b_0$ sharpened by the term u^2 of (7) — is what reaches $l \le 4 \cdot 10^4$ with no bound on k. The frontier is now $l = 10^5$, and it is left open with a measured price (Remark 9.5): about fifty-two minutes of processor time for a single evaluation, and a level-1 discriminant that leaves the evaluator's machine-integer range near $l \approx 6.6 \cdot 10^4$, so that the next decade needs either a compiled kernel or a different arithmetic. What no amount of that would settle is Remark 5.3 itself, which is a statement about all l; the two sub-families displayed above are where a proof would have to bite. And the growth tables of [6, Remark 5.4], for the families $(1, 1 - kx^2)$, $(1, 1 + kx^3)$ and $(1, 1 + kx^k)$, are a separate problem: of the first of them [6, Example 7.1] records that the authors' own implementation of the procedure that determines m_0 "gets quite slow" beyond $k = 24$, and our own estimate of the cost of reproducing the tables is tens of CPU-hours.

Changes from version 1 of this paper. Nothing stated in version 1 is withdrawn: every theorem, lemma, corollary and proposition of that version stands here with its statement unchanged, and the earlier computation it rests on is unchanged as well. What is added is Section 8 (the chord bound and the two levels of the search), Section 9 (Theorem 9.1, Propositions 9.2 and 9.3, and two labelled computations), Proposition 10.3, and items (5) and (6) of Section 1; the abstract now states the wider range and the two exact values. Four passages of version 1 are updated rather than left standing, and we say which. The sentence closing Section 10 said that the exact values 4 and 3 were reproduced numerically but certified only in the form ≥ 3; Proposition 10.3 now certifies them. The paragraph closing this section listed three continuations, of which two are now carried out, and it names the new frontier in place of the old one. The opening of this section, and the discussion of the surviving region that follows Remark 11.1, now record what Theorem 9.1 settles and what it leaves. And Section 12 is rewritten to cover the added material. Nothing in version 1 was found to be wrong.

12. Verification

Machine-checked in Lean 4 [2] against Mathlib, in the form stated, are Theorems 3.1, 5.1, 5.2, 7.1 and 9.1, Corollaries 3.2 and 7.2, Lemmas 4.1, 4.2, 6.1 and 6.2, and Propositions 9.2, 9.3, 10.1, 10.2 and 10.3 — of Lemma 4.2 in the form of its last sentence, which is the implication every later result uses, the equivalence with (2)–(5) being unwound inside its formal proof, and of Corollary 3.2 in the direction that carries its content, that stability of the single set T_{M+1} gives $m_0(c, d) \le M$; its converse is the definition of $m_0(c, d) \le M$ read at $m = M + 1$. Lemma 8.1 is checked in the form in which the development consumes it: the Lean statement is that a passing verdict of the walk over the window described after it transfers to every integer solution of the level, and the two inequalities of (7) are steps of that proof rather than declarations of their own. The weight $\|\cdot\|$, the truncation $f|_m$, the sets T_m, stability, the relation $m_0(c, d) \le M$ and weak robustness are defined in Lean directly from the sentences of [6] quoted in Sections 1 and 2; Mathlib supplies the polynomial algebra and nothing else, having no notion of polynomial weight. Everything in Sections 3–6 and Section 8, together with Propositions 9.3 (ii), 10.1 and 10.2, is checked by the Lean kernel alone, depending on no axioms beyond propositional extensionality, quotient soundness and choice; in particular Theorem 5.2, the infinite sub-family, is free of any evaluation step, and so are the reductions of T_4 and T_5 to their systems and the witness $a = b = 2 + 2x - x^2 + x^3$ used in Proposition 10.3. The development contains eight finite evaluations in all, each carried out by Lean's compiled evaluator (`native_decide`) and each contributing one evaluation axiom: the enumeration described in the

proof sketch of Theorem 7.1, on which that theorem and Corollary 7.2 depend; the enumeration described in the proof sketch of Theorem 9.1, on which that theorem and Proposition 9.2 depend; the two searches that decide $m_0(1+3x, 9) \le 3$ and $m_0(1+2x, 4) \le 4$ in Proposition 10.3; the three false verdicts of Proposition 9.3 (i); and one further control, the false verdict of the depth-4 search at $(2, 4)$, whose mathematical counterpart — that 4 is not stable there — is the kernel-checked witness just named. The soundness of every one of these enumerations — that a *true* verdict implies the non-existence of integer solutions, including the divisibility bookkeeping by which the b_i are recovered, the bound $x^2 \le n \Rightarrow |x| \le \lfloor\sqrt{n}\rfloor$ used to cut the ranges, the windows of Lemma 8.1, and the derivation of the shortened divisor test from weak robustness — is itself proved in the kernel; the searches allocate nothing, each loop being an index walk. There is no `sorry` in the development and it declares no axiom of its own. Not formalised, and labelled as such where they appear, are Remarks 2.3, 5.3, 6.3, 9.4, 9.5, 11.1 and 11.2: all but Remark 11.1 are computations and measurements carried out by independently written C programs, and Remark 11.1 is an ordinary paper proof which we have not formalised because it needs the greatest-common-divisor parametrisation of the solution lattice. Remark 6.4's two witnesses are verified by the substitutions displayed there. The repository is private; repository reference to be added at submission.

References

[1] M. Filaseta, K. Ford and S. Konyagin, *On an irreducibility theorem of A. Schinzel associated with coverings of the integers*, Illinois Journal of Mathematics **44** (2000), no. 3, 633–643; doi:10.1215/ijm/1256060421.

[2] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, in: A. Platzer and G. Sutcliffe (eds.), Automated Deduction — CADE 28, Lecture Notes in Computer Science **12699**, Springer, Cham, 2021, pp. 625–635; doi:10.1007/978-3-030-79876-5_37.

[3] W. Ljunggren, *On the irreducibility of certain trinomials and quadrinomials*, Mathematica Scandinavica **8** (1960), 65–70; doi:10.7146/math.scand.a-10593. The 1960 argument whose trick [6] generalises; cited from [6] and not consulted here.

[4] A. Schinzel, *Reducibility of lacunary polynomials. I*, Acta Arithmetica **16** (1969), 123–159; doi:10.4064/aa-16-2-123-160. (The publisher's index and Crossref record the range as 123–160; the article itself ends on p. 159.)

[5] A. Schinzel, *Polynomials with special regard to reducibility*, Encyclopedia of Mathematics and its Applications **77**, Cambridge University Press, Cambridge, 2000, with an appendix by U. Zannier; doi:10.1017/CBO9780511542916. [6, Theorem 1.2] is deduced from its Theorem 74; cited from [6] and not consulted here.

[6] W. Sawin, M. Shusterman and M. Stoll, *Irreducibility of polynomials with a large gap*, Acta Arithmetica **192** (2020), no. 2, 111–139; doi:10.4064/aa180526-12-6; arXiv:1803.10811v1. All quotations above are from the e-print, whose arXiv listing carries the single version v1 of 28 March 2018, and every numbered cross-reference above is to that version's numbering, read off the labels its own LaTeX source resolves: Theorem 1.2, display (4.1), Definition 4.2, Corollary 4.4, Lemma 4.5, Lemmas 5.1 and 5.2, Remarks 5.3 and 5.4, and Examples 7.1 and 7.5. The printed Acta Arithmetica version is behind the publisher's paywall and was not consulted; its numbering may differ.