Spectral bounds for the number of Robin candidates

Circulation draft - revised September 15, 2026.

Prepared for specialist mathematical and attribution review. The main written theorem is stated relative to named classical imports; the companion Lean verification has the narrower scope recorded in Section 8.

Abstract

We give a quantitative upper bound for the number of integers satisfying the reverse weak Robin inequality. The argument combines a weighted zero sum, an exact centered Mertens identity and a count of low-benefit prime-support patterns. Under an explicit zero-free-region and uniform zero-density contract, the resulting rate is expressed through the exact infimum W_2(y)=inf_t[eta(t)log y+2log t] and a specified secondary factor. Classical and Vinogradov-Korobov inputs give unconditional ordinary-zeta specializations. The proof is self-contained relative to the classical imports stated below. Its relationship to earlier counting and integrated Mertens results is still under investigation; no priority claim is made.

1. Main theorem and inputs

All logarithms are natural. Primes are ordinary positive rational primes; phi is Euler's totient, sigma is the sum-of-divisors function, and gamma_E is the Euler-Mascheroni constant. We count the weak reverse inequality, including equality, for n>=3. The finitely many small integers are retained throughout. The asymptotic symbols refer to the indicated real variable tending to infinity; constants may depend on fixed analytic inputs unless stated absolute.

Let eta be a continuous, positive, nonincreasing zero-free function for ordinary zeta above a fixed T>*1, with eta<=1/2 and eta(t)->0. Thus every nontrivial zero rho=beta+i gamma with |gamma|>=T* obeys beta<=1-eta(|gamma|). Admit a zero-density bound, with fixed A_d>0, B>=1, C>=0, C_d>0, sigma_0 in (1/2,1),

```
N(sigma,T)<=C_d T^[A_d(1-sigma)^B](log T)^C
```

uniformly for sigma_0<=sigma<1 and T>=T. *Here N(sigma,T) counts zeros with beta>sigma and 0<gamma<T, with multiplicity; conjugate zeros are included separately when summing. Endpoint conventions can be replaced by one-sided limits in the integrals. Enlarge T* once to meet the zero-free and density ranges.

Write, for y sufficiently large and u=log y,

```
W(y)=W_2(y)=inf_{t>=T*}[eta(t)u+2log t],
P_0(y)=(u/W(y))^(BC) exp(A_d W(y)^(B+1)/u^B),
g(y)=y sqrt(P_0(y))exp(-W(y)/2)/u.
```

Assume W=o(u) and W/log u->infinity. Retain these exact definitions. Admit the usual zeta facts 0<beta<1, no real nontrivial zeros, the total zero count N(T)=O(T log T), the once-integrated explicit formula specified in Section 3.2, and the classical Mertens normalization. For the combinatorial step admit ordinary-prime Brun-Titchmarsh and ENS Lemmas 1-2 in the precise form stated in Section 2.3.

Define for real X>=3

```
R(X)=#{n integer : 3<=n<=X,
       sigma(n)>=exp(gamma_E)n log log n}.
```

Theorem 1. Under these inputs there exist fixed L>0 and X_0 such that

```
R(X)<=exp(L g(log X))  for every X>=X_0.            (T1)
```

L and X_0 may depend on the fixed analytic inputs. No uniform numerical L or certified numerical X_0 is supplied. When the inputs are instantiated by their classical unconditional versions for ordinary zeta/primes, this is an unconditional written consequence of those theorems. The arbitrary eta version remains explicitly conditional on the displayed input contract.

The additional imported statements are the primitive explicit formula in Section 3.2, the Mertens limit log M(y)-log log y -> gamma_E in Section 4.2, and the prime-counting, prime-index and partition estimates in Section 2.3. Elementary Chebyshev and prime-power bounds are proved or reduced to finite sums below. References and the depth of the source checks are recorded in Section 7.

The proof has four separate steps: count a benefit relaxation for a positive budget H; construct H by spectral estimates; place all candidates in that relaxation by an exact identity; and aggregate full unit shells. The counting lemma alone constructs no H and says nothing about candidate inclusion. Neither inclusion nor counting uses a local theta-window error.

2. Counting a low-benefit relaxation

2.1. Definitions and counting proposition

Let y tend to infinity, u=log y, and let H(y)>0 eventually satisfy

```
log H(y) / log y -> 0,        H(y)(log y)^3 -> 0.          (P1)
```

Use r(t)=log(t/(t-1)), lambda=r(y)/u, M(y)=product_{p<=y}(1-1/p)^(-1), theta(y)=sum_{p<=y} log p, and the totient benefit

```
B_phi(n)=log M(y)-log(n/phi(n))+lambda(log n-theta(y)).
```

Let E_y=E_y(H) be the positive integers in the unit log shell [y-1,y] with B_phi(n)<=H(y).

Proposition 2. There is an absolute K_count>0 and an eventual threshold, allowed to depend on H, such that

```
|E_y| <= exp(K_count G(y)),        G(y)=y sqrt(H(y)/u).          (P2)
```

The same bound applies to the sigma-benefit shell of Section 4.1, by B_sigma>=B_phi. This is an upper count of a relaxation. Section 4 supplies the separate candidate inclusion.

2.2. Repeated powers and distant support

Write n=s*m with s=rad(n). The nonnegative local costs imply

```
B_phi(n)=B_phi(s)+lambda log m.
```

Since r(t)/log t is strictly decreasing for t>1, all support costs are nonnegative; further copies cost lambda log p each. Thus the number of possibilities for m is at most exp(H/lambda). Omitted primes p<y/2 each cost at least 1/(2y); their subsets have at most (floor y+1)^(floor(2yH)+1) possibilities. For included p>2y,

```
r(p)<=1/(p-1)<1/(2y-1)<=(2/3)lambda log p,
```

so the product of these primes is at most exp(3H/lambda). Integer size alone bounds the number of products by the same expression. Together these give an outside-choice count at most exp(F_out), where

```
F_out=4yuH+(2yH+1)log(y+1).                                  (P3)
```

For completeness, the omitted-prime estimate follows by setting f(t)=r(t)-lambda log t. Then f(y)=0 and

```
-f'(t)=1/[t(t-1)]+lambda/t >= 1/y^2   (2<=t<=y).
```

Integrating over [y/2,y] gives 1/(2y). Also 1/lambda<=yu, as used in (P3). These arguments require only y>=4. They count enlargements of the permitted data; no claim of independent compatibility is needed for an upper bound.

2.3. Near support and imported prime estimates

All remaining choices are bits: whether a reference prime in [y/2,y] is omitted, or a prime in (y,2y] is included. Set

```
d(p)=|r(p)-lambda log p|.
```

The derivative in Section 2.2 yields, on the entire interval [y/2,2y],

```
d(p) >= |p-y|/(4y^2).                                    (P4)
```

Use the disjoint lower interval [y/2,y] and upper interval (y,2y]. List each side's primes in order of increasing distance from y, calling them P_i with one-based indices, and put h_i=|P_i-y|. Free at most the first two bits on each side, paying a factor at most 16. A prime equal to y occurs only in the lower list and has distance zero; it is among these free bits.

For i>=3 put j=i-1>=2. On the upper side, (y,P_i] contains i of the listed primes. On the lower side, (P_i,y] contains i-1, excluding P_i itself. Thus each interval contains at least j primes. Distinct integer locations also give h_i>=j>=2. The imported Brun-Titchmarsh bound is pi(x+h)-pi(x)<=2h/log h for x>=0 and h>=2 [ENS, K]. The two applications have base points y and P_i respectively, both positive. Therefore

```
j <= 2h_i/log h_i,
h_i >= (j log j)/2 >= 0.23 q_j >= (1/5)q_j,              (P5)
```

where q_j is the j-th ordinary prime, q_1=2 and q_2=3. The penultimate inequality is ENS Lemma 1 [ENS], j log j>=0.46 q_j for j>=2. The argument is the same for the upper and lower lists; using j=i-1 deliberately permits the lost endpoint in the lower list. The interval lengths here are at least 2, so no log-zero case is used.

For each side, let b_i be its changed-support bits after the first two. Since all costs are nonnegative, (P4)-(P5) imply

```
sum_{i>=3} q_{i-1} b_i <= 20 y^2 H = T.                  (P6)
```

Enlarge from bits to arbitrary nonnegative integer coefficients, and from the finite list to all ordinary primes. Let Pi(T) count all such prime partitions of total weight at most T, equivalently the finite-support sequences (x_j) of nonnegative integers with sum_j q_j x_j<=T. This is a finite family for T>=0: an occupied prime has q_j<=T and its coefficient is at most floor(T/q_j). The constrained bits inject into this family by assigning b_i to the distinct prime q_(i-1); the free bits retain the omitted coordinates. Hence the finite near-pattern count is at most 16 Pi(T)^2, before any asymptotic estimate is applied. Separately, ENS Lemma 2 [ENS] supplies

```
log Pi(T) = O(sqrt(T/log T)),       T -> infinity.       (P7)
```

There is no dependence on the budget restriction in ENS Theorem 6: (P7) is a one-variable partition estimate, and the uniform interval estimate (P5) was applied directly to the actual primes. Both sides may use the whole budget H; that only enlarges the count.

2.4. Completion of the counting proof

The outside data, near bits, and m determine n uniquely: the omitted and included support slots determine rad(n), while m records all additional prime multiplicities. The slots for first presence and further copies are disjoint. Allowing incompatible encoded data only enlarges this upper count. Thus the explicit intermediate estimate is

```
|E_y| <= 16 exp(F_out) Pi(20y^2H)^2.                     (P8)
```

Under (P1), T=y^(2+o(1)) tends to infinity and log T~2u. Therefore sqrt(T/log T)=O(G). Moreover

```
yuH/G=sqrt(Hu^3) -> 0,
log(y+1)/G -> 0,
```

the latter since G=y^(1+o(1)). Hence F_out=o(G), proving (P2) as written. In fact the encoding never uses shell membership, so it also bounds the full totient-benefit sublevel; the stated shell conclusion needs no such extension.

The partition input has only the variable T. In particular, its implied constant does not acquire hidden y-dependence when T=20y^2H(y). The full ENS integer-counting theorem is not invoked outside its printed budget range. The leading asymptotic for log Pi(T) is (2pi/sqrt(3))sqrt(T/log T); we use only its O bound and do not specify K_count numerically.

3. Constructing the spectral budget

A standalone analytic lemma records this step's exact input contract, constant dependencies and specialist-review questions. It is a separate presentation of the proof below, not a stronger result.

Define theta(t)=sum_(p<=t)log p and psi(t)=sum_(p^k<=t)log p. Let rho=beta+i gamma run over all nontrivial zeros with multiplicity, counting both signs of gamma. Put

```
S(y)=sum_rho y^(beta-1)/|rho(rho-1)|,
M_z=sum_rho 1/|rho(rho-1)| < infinity.
```

The total zero count implies convergence of M_z; the finitely many bounded height zeros have nonzero denominators. Each summand of S is decreasing in y. For q>0 write W_q(y)=inf_(t>=T*)[eta(t)u+q log t], with u=log y.

3.1. A split uniform in the moving parameter

For sufficiently small 0<r<1, set

```
delta=(r/(2A_d))^(1/B),    sigma_r=1-delta.
```

Then sigma_r>=sigma_0 and the density exponent at sigma_r equals r/2. Stieltjes integration, counting both signs of gamma, gives

```
sum_{beta>sigma_r, |gamma|>=T*} |gamma|^(-r)
  <= 2r integral_{T*}^infinity N(sigma_r,t)t^(-r-1)dt
  <= 2C_d r integral_1^infinity t^(-1-r/2)(log t)^C dt
  = 2^(C+2) C_d Gamma(C+1) r^(-C).                    (U1)
```

The lower boundary contribution is nonpositive and can be dropped; the upper boundary vanishes since N(sigma_r,T)T^(-r)=O(T^(-r/2)(log T)^C). Endpoint atoms can instead be handled with left limits, with the same bound. The gamma integral is valid for every real C>=0, not just integer C. The constant in this estimate is independent of r; write K1=2^(C+2) C_d Gamma(C+1).

Finitely many zeros below T* lie to the left of sigma_r for all sufficiently small r. For every zero with beta<=sigma_r, including these finite low zeros,

```
sum y^(beta-1)/|rho(rho-1)| <= M_z exp(-delta u),
M_z=sum_rho 1/|rho(rho-1)|<infinity.
```

For the remaining zeros, |rho(rho-1)|>=gamma^2, and the zero-free region gives

```
y^(beta-1)/|rho(rho-1)|
  <=exp(-[(1-beta)u+(2-r)log|gamma|]) |gamma|^(-r)
  <=exp(-W_{2-r}(y)) |gamma|^(-r).
```

Sum this inequality using (U1). It proves

```
S(y) <= M_z exp(-delta u)+K1 r^(-C)exp(-W_{2-r}(y)).  (U2)
```

There is no substitution into a theorem with uncontrolled constants: (U2) is proved uniformly first.

Endpoint substitution.

Set

```
delta=W/u,    r=2A_d(W/u)^B.
```

Since W=o(u), all preceding small-r conditions hold eventually. Moreover, pointwise comparison inside the infimum gives

```
W_{2-r}(y) >= (1-r/2) W(y).
```

Indeed eta(t)u+(2-r)log t is at least (1-r/2) times eta(t)u+2log t. Substitute this and delta u=W into (U2):

```
S(y) << P_0(y) exp(-W(y)),                              (U3)
P_0(y)=(u/W(y))^(BC) exp(A_d W(y)^(B+1)/u^B).
```

The fixed factor (2A_d)^(-C) is absorbed into the constant. Eventually u/W>1, so P_0>=1 and also absorbs the first term of (U2). In particular

```
log P_0=BC log(u/W)+A_d W^(B+1)/u^B=o(W).
```

The result reaches the endpoint with a specified secondary loss. It does not in general show S=O(exp(-W)).

3.2. The primitive formula and the integrated tail

We import the classical once-integrated explicit formula [K], normalized as

```
F(x)=integral_2^x (psi(t)-t)dt
    =-sum_rho x^(rho+1)/(rho(rho+1))-x log(2pi)+O(log x).
```

The zero series is absolutely and uniformly convergent on compact intervals: the numerators are bounded there and the reciprocal quadratic denominators are summable. On a finite interval [y,T], integration by parts followed by termwise integration therefore gives

```
integral_y^T (psi(t)-t)/t^2 dt
  =sum_rho [y^(rho-1)-T^(rho-1)]/[rho(rho-1)]+O(1/y).
```

The error is uniform in T>=y>=2. The linear term contributes O(1/y), and the primitive O(log x) remainder contributes O(log y/y^2). Since beta<1, the T-dependent zero sum tends to zero by dominated convergence, dominated by the summable reciprocal denominators. Thus the improper integral exists and

```
J_psi(y)=integral_y^infinity (psi(t)-t)/t^2 dt
       =sum_rho y^(rho-1)/(rho(rho-1))+O(1/y),
|J_psi(y)|<=S(y)+C_J/y.
```

This avoids interchanging an unsmoothed zero series with an infinite integral. It also requires no reciprocal factor (1-beta)^(-1).

3.3. The weighted theta tail and the prime-power correction

Let Delta(t)=theta(t)-t, and define I_psi and I_theta by integrating, respectively, psi(t)-t and Delta(t) against a(t)/t^2 below. Their existence will follow from the following argument.

The function S(y)+C_J/y is a decreasing exact envelope: each zero has Re(rho)<1 and the nonnegative series converges. Set

```
a(t)=1/log t+1/log^2 t.
```

Integration by parts first on a finite interval, then at infinity, gives

```
I_psi(y)=a(y)J_psi(y)+integral_y^infinity a'(t)J_psi(t) dt,
|I_psi(y)|<=2a(y)(S(y)+C_J/y).                        (I8)
```

The endpoint a(T)J_psi(T) vanishes, and the integral against |a'| converges by the decreasing envelope and integral_y^infinity |a'|=a(y). This establishes existence of the weighted improper integral without requiring the majorant P_0(t)exp(-W(t)) itself to be decreasing.

Let J_theta be the integral defining J_psi with psi replaced by theta. Its convergence follows from the first correction below. Keep the two prime-power corrections distinct:

```
Q_J(y)=J_psi(y)-J_theta(y)
      =integral_y^infinity (psi(t)-theta(t))/t^2 dt >=0,
Q_I(y)=I_psi(y)-I_theta(y)
      =integral_y^infinity (psi(t)-theta(t))a(t)/t^2 dt >=0.
```

There are at most log(t)/log(2) exponents k>=2 in psi-theta. Each inner sum over p<=t^(1/k) is at most sqrt(t)log t, by counting integers instead of primes. Hence psi(t)-theta(t)=O(sqrt(t)log^2 t). This yields a convergent Q_J=O(log^2 y/sqrt y), and, using a(t)<=2/log t for t>=4,

```
Q_I=O(log y/sqrt y),    I_theta=I_psi-Q_I.            (I9)
```

For example, integral_y^infinity (log t)t^(-3/2)dt =(2log y+4)/sqrt y. This proves existence of I_theta and controls its absolute value by |I_psi|+Q_I. No O(1/y) estimate for Q_J is used. A sharper Theta(1/sqrt y) statement for Q_J is separate and unnecessary here.

Now set A(y)=P_0(y)exp(-W(y))/u. Then

```
(log y/sqrt y)/A(y)=u^2 exp(W-u/2)/P_0 -> 0,
(1/(yu))/A(y)=exp(W-u)/P_0 -> 0.
```

Together with (I8), these yield |I_theta|<=C_I A eventually for a fixed C_I after absorbing the prime-power and explicit-formula remainders. The fixed C_I includes these analytic remainders. The later choice of K_0 will leave only the two arithmetic corrections in Section 4 to absorb.

4. Centered inclusion of actual candidates

4.1. Benefits and the exact threshold

Retain y>=4 and u=log y. For clarity, the benefit definitions are

```
theta(y)=sum_{p<=y} log p,    Delta(y)=theta(y)-y,
r(t)=log(t/(t-1)),           lambda=r(y)/u,
M(y)=product_{p<=y}(1-1/p)^(-1),
B_phi(n)=log M(y)-log(n/phi(n))+lambda(log n-theta(y)),
B_sigma(n)=log M(y)-log(sigma(n)/n)+lambda(log n-theta(y)).
```

Use the candidate convention

```
C_y={n integer : n>=3, y-1<=log n<=y,
                sigma(n)>=exp(gamma_E) n log log n}.
```

This includes equality at the threshold and the finite small-n range. It does not assert that any candidate exists. Define the full relaxation

```
E_y(H)={n integer : n>=1, y-1<=log n<=y, B_phi(n)<=H}.
```

For ell in [y-1,y], put

```
U_y(ell)=log M(y)-gamma_E-log(log ell)+lambda(ell-theta(y)).
```

Here ell>=3, so every logarithm is defined and log ell>0. For n>=3 in the shell, taking logarithms of the positive Robin threshold gives exactly

```
n in C_y  iff  B_sigma(n)<=U_y(log n).              (I1)
```

The Euler product for sigma gives sigma(n)/n<=n/phi(n); hence B_phi(n)<=B_sigma(n). Both benefits are nonnegative: for B_phi, omission below y costs r(p)-lambda log p>=0, first inclusion above y costs lambda log p-r(p)>=0, and each additional copy costs lambda log p>=0. These signs follow from the decrease of r(t)/log t for t>1. Then use B_sigma>=B_phi. Therefore a uniform upper bound U_y<=H implies

```
C_y subset {shell, B_sigma<=H} subset E_y(H).       (I2)
```

The first inclusion uses (I1), not merely the inequality between benefits.

4.2. Exact centered identity

Admit the classical Mertens normalization

```
log M(T)-log log T -> gamma_E.
```

The source normalization is recorded in [Z]. Section 3 proves existence of the following improper integral under the theorem inputs:

```
I_theta(y)=integral_y^infinity Delta(t)(log t+1)/(t^2 log^2 t) dt.
```

Put

```
R_tail(y)=sum_{p>y} sum_{j>=2} 1/(j p^j) >= 0,
S_1(y)=sum_{p<=y} 1/p,    Q_1(y)=sum_{p<=y} sum_{j>=2} 1/(j p^j),
k(t)=1/(t log t).
```

The total Q_1(infinity) is finite by comparison with sum_{m>=2} 1/(m(m-1)). Partial summation on (y,T], with theta counting primes <= its argument, gives the exact finite-interval identity

```
S_1(T)-S_1(y)-log log T+log log y
  =Delta(T)k(T)-Delta(y)k(y)
   +integral_y^T Delta(t)(log t+1)/(t^2 log^2 t) dt.
```

There is no half-weight convention at a prime endpoint. The atom at y is included in theta(y) and excluded from (y,T]; changing a single endpoint of the ordinary integral does not change its value.

For completeness, an elementary Chebyshev bound suffices for the boundary limit. For integers m>=1, the product of primes in (m,2m] divides binomial(2m,m), so theta(2m)-theta(m)<=2m log 2. Summing at powers of two and then enclosing a real t in the next power of two gives

```
0<=theta(t)<=4(log 2)t,    |Delta(t)|<=(4log 2+1)t  (t>=2).
```

Thus Delta(T)k(T)->0. Use log M=S_1+Q_1 and the Mertens limit above to pass to T=infinity, obtaining

```
log M(y)-gamma_E-log log y
  =Delta(y)/(yu)-I_theta(y)-R_tail(y).
```

Substitution into U_y(y) yields

```
U_y(y)=(1/(yu)-lambda)Delta(y)-I_theta(y)-R_tail(y).      (I3)
```

The -R_tail sign is useful for an upper bound. No sign of I_theta is assumed.

4.3. Uniform arithmetic corrections

The convergent logarithm series, with y>=4, gives

```
0<=r(y)-1/y=sum_{k>=2} 1/(k y^k)<=1/(2y(y-1)),
0<=lambda-1/(yu)<=1/(y^2u).                          (I4)
```

Consequently the absolute endpoint contribution in (I3) is at most c_theta/(yu), where c_theta=4log 2+1.

For t in [y-1,y], t>=y/2, log t>=u/2, and log t>=log 3>1. Hence

```
|k'(t)|=(log t+1)/(t^2 log^2 t)<=16/(y^2u).
```

Differentiate the full threshold, keeping y fixed:

```
U_y'(ell)=lambda-k(ell)
         =[lambda-k(y)]+[k(y)-k(ell)].
```

By (I4) and the mean value bound for k, uniformly on the whole closed shell,

```
|U_y'(ell)|<=17/(y^2u),
|U_y(ell)-U_y(y)|<=17/(y^2u).                        (I5)
```

Thus, including endpoints and allowing either sign of Delta and I_theta,

```
sup_{ell in [y-1,y]} U_y(ell)
  <=max(0,-I_theta(y))+c_theta/(yu)+17/(y^2u).       (I6)
```

These explicit arithmetic constants are convenient majorants; they do not supply an explicit constant for the imported spectral bound or final count.

4.4. Positive-budget inclusion

Proposition 3. Let A(y)>0 eventually, yA(y)->infinity, and suppose for a fixed C_I>=0 that

```
max(0,-I_theta(y))<=C_I A(y)                         (I7)
```

for all sufficiently large y. Retain the normalization and integral existence conditions of Section 4.2. Choose any fixed K_0>C_I and set H(y)=K_0 A(y). Then eventually sup U_y<=H(y), and (I2) holds.

Proof. The two correction terms in (I6), divided by A, tend to zero:

```
c_theta/(yuA)+17/(y^2uA) -> 0.
```

Set epsilon=K_0-C_I>0. Beyond one common threshold they sum to at most epsilon A, and (I6)-(I7) give sup U_y<=K_0 A. Combine with (I1)-(I2). All constants are fixed independently of y. The threshold can depend on the inputs and the margin K_0-C_I. QED.

The sufficient growth assumption yA->infinity is deliberately generous, not claimed necessary. Likewise (I7) is a sufficient one-sided input; the stronger absolute estimate |I_theta|<=C_I A also supplies it. No extra signed cancellation is needed once an absolute bound is available at the smaller spectral scale.

If sup U_y<0, nonnegative B_sigma already makes C_y empty. If sup U_y=0, no emptiness assertion is needed: the same strictly positive H works. Never substitute the possibly zero positive-part allowance into (P1).

4.5. The single-shell consequence

Take A(y)=P_0(y)exp(-W(y))/u from Section 3 and fix K_0>C_I. Set H=K_0 A. The exact secondary factor satisfies log P_0=o(W), and the theorem assumes W=o(u), W/log u->infinity. Therefore

```
log(yA)=u+log P_0-W-log u -> infinity,
log H=log K_0+log P_0-W-log u=o(u),
log(Hu^3)=log K_0+log P_0-W+2log u -> -infinity.
```

The same positive H both bounds the full candidate threshold, by Proposition 3, and satisfies (P1). Proposition 2 consequently gives

```
|C_y|<=|E_y(H)|<=exp(K_count y sqrt(H/u))
     =exp(K_count sqrt(K_0) g(y)).                    (SHELL)
```

K_count is absolute; K_0 depends on the fixed analytic inputs. Choose one threshold meeting all preceding eventual requirements. The construction of H, the inclusion and the count are separate deductions, and no local theta-error hypothesis has entered them.

5. Exact-rate regularity and global aggregation

5.1. A finite shell cover

Let C_y consist of integers n>=3 with y-1<=log n<=y and sigma(n)>=exp(gamma_E)n log log n. Let R(X) count the same threshold candidates with 3<=n<=X, for real X>=3.

Suppose g is positive and nondecreasing eventually, log y=o(g(y)), and there is a fixed D_shift>=1 such that

```
g(y+h)<=D_shift g(y)  (0<=h<=1, y sufficiently large).
```

Suppose also there are fixed L>0 and y_0 such that

```
|C_y|<=exp(Lg(y))  for every real y>=y_0.            (G1)
```

It suffices to assume (G1) at integer y. Then there is a fixed L_global>0 such that eventually

```
R(X)<=exp(L_global g(log X)).                        (G2)
```

Proof. Fix an integer J>=4 beyond all eventual thresholds. Put m=ceil(log X) and B_0=ceil(exp J), a fixed integer. A candidate with log n<=J lies in the initial interval counted by B_0. Every remaining candidate belongs to C_j with j=ceil(log n), where J<j<=m. At an integer logarithm the endpoint convention is still valid; overlapping shell endpoints only enlarge an upper count. Hence, for m>=J,

```
R(X)<=B_0+sum_{j=J+1}^m |C_j|
     <=B_0+(m+1)exp(Lg(m))
     <=(B_0+m+1)exp(Lg(m)).                          (G3)
```

The last step uses exp(Lg(m))>=1. With y=log X, m-y in [0,1), so

```
log[(B_0+m+1)exp(Lg(m))]
  <=L D_shift g(y)+log(B_0+m+1).
```

Since m/y->1 and log y=o(g(y)), the final term is o(g(y)). For example L_global=L D_shift+1 works after enlarging the threshold. This bounds R(X) without taking log of R(X), which may be zero. QED.

The number of shells is O(log X); the quantity absorbed in the exponent is its logarithm O(log log X), not the number of shells itself. A finite initial segment never needs to satisfy the asymptotic shell hypothesis.

5.2. Regularity of the exact infimum

To separate variables, write

```
w(u)=W_2(exp u)=inf_{t>=T*} [eta(t)u+2log t],  T*>1.
```

Assume 0<eta(t)<=1/2, with the established spectral limits

```
w(u)=o(u),    w(u)/log u -> infinity.                (G4)
```

For the actual spectral secondary factor, with fixed A_d>0, B>=1, C>=0, define

```
p(u)=(u/w(u))^(BC) exp(A_d w(u)^(B+1)/u^B),
g(exp u)=exp(u) sqrt(p(u)) exp(-w(u)/2)/u.            (G5)
```

All quantities are eventually positive, and p>=1 eventually. These are exact definitions, not asymptotic replacements inside an exponential.

For 0<u<=v the affine functions defining w have nonnegative slopes and positive intercepts. Taking infima gives

```
w(u)<=w(v)<=w(u)+(v-u)/2,
w(v)<=(v/u)w(u).                                      (G6)
```

For the second inequality, each eta(t)v+2log t is at most (v/u)[eta(t)u+2log t]; choose arbitrarily close minimizers if necessary. Thus w is locally Lipschitz, hence locally absolutely continuous, and at almost every u>0,

```
0<=w'(u)<=w(u)/u.                                     (G7)
```

Indeed the right difference quotients are bounded by w(u)/u using (G6). No unique minimizer, derivative of eta, or everywhere differentiability of w is assumed. One could also note concavity of w, but it is not needed.

5.3. Regularity of the secondary factor

Put b=BC and q=w/u. On each sufficiently large compact interval, w is positive; compositions in (G5) are locally absolutely continuous. At almost every u,

```
(log p)'=b(1/u-w'/w)
          +A_d q^B[(B+1)w'-Bq].                       (G8)
```

By (G7), 0<=w'<=q. Therefore

```
0<=b(1/u-w'/w)<=b/u,
-A_d B q^(B+1)
  <=A_d q^B[(B+1)w'-Bq]<=A_d q^(B+1).
```

In particular

```
|(log p)'|<=b/u+A_d B q^(B+1) -> 0.
```

Here the last limit follows from q=w/u->0 and B>=1. Writing L_rate(u)=log(g(exp u)), (G5) gives

```
L_rate'=1-1/u+(log p)'/2-w'/2.
```

Its distance from 1 is bounded almost everywhere by the deterministic quantity

```
1/u+b/(2u)+(A_d B/2)q^(B+1)+q/2 -> 0.
```

Choose u_0 beyond which this bound is at most 1/2. The fundamental theorem for absolutely continuous functions then gives, for u_0<=u<=v,

```
(v-u)/2 <= L_rate(v)-L_rate(u) <= 3(v-u)/2.           (G9)
```

This justifies the monotonicity inference even at nondifferentiable points. In particular g is eventually increasing, and for y>=exp u_0, h in [0,1],

```
1<=g(y+h)/g(y)<=((y+h)/y)^(3/2).
```

One may use the fixed shift bound D_shift=3 for y>=1.

Also

```
log p=b log(u/w)+A_d w(w/u)^B=o(w)=o(u).
```

The first term is O(log u)=o(w), since eventually 1<w<u. Thus log(g(exp u))=u+o(u), and g(y)=y^(1+o(1)); in particular log y=o(g(y)). All hypotheses of Section 5.1 are now proved for the exact factor (G5).

Scope. This proof covers the stated spectral infimum and specific secondary factor under (G4), not only smooth classical/VK shapes. Merely assuming log P_0=o(W) for an arbitrary positive P_0 would not justify monotonicity or bounded shifts. Such a general budget needs separate rate regularity or a maximum-envelope version of (G3).

Completion of Theorem 1. Apply Section 5.1 to (SHELL) with L=K_count sqrt(K_0). Sections 5.2-5.3 establish positivity, monotonicity, bounded unit shifts and log y=o(g(y)) for the exact theorem rate. Enlarge the fixed shell threshold once to include all analytic and counting thresholds. The finite initial segment is absorbed in (G3), and (T1) follows. No logarithm of an empty candidate count is taken.

6. Classical and Vinogradov-Korobov corollaries

For an admitted classical zero-free region eta(t)=a/log t eventually, a>0, the minimizer is eventually interior and elementary minimization gives

```
w(u)=2sqrt(2au).
```

With the log-free density parameters B=1,C=0, p=exp(8A_d a) is a fixed constant. Consequently there are fixed L,X_0 such that R(X)<=exp(L g(log X)) for X>=X_0, where the global rate is

```
g(log X) = constant * log X/log log X
             *exp(-sqrt(2a log log X)).              (G10)
```

For an admitted VK region eta(exp v)=a v^(-2/3)(log v)^(-1/3) eventually, retain the exact infimum w(u). Only for bounding the secondary factor use

```
w(u) comparable to u^(3/5)/(log u)^(1/5).
```

Here is an order proof sufficient for that use. Let v_0=u^(3/5)/(log u)^(1/5). Evaluating at v_0 bounds w above by a constant times v_0. For v>=v_0 the term 2v gives a lower bound 2v_0. For the remaining sufficiently large v<=v_0, the decreasing function v^(-2/3)(log v)^(-1/3) is at least its value at v_0, giving a positive constant times v_0. The fixed initial v interval has eta bounded below by a positive constant and eventually contributes a much larger multiple of u. Thus w is comparable to v_0. This also verifies (G4).

With B=3/2,C=15,

```
w^(5/2)/u^(3/2)=O((log u)^(-1/2)),
sqrt(p(u))=O((u/w(u))^(45/4)).
```

Consequently (G2) implies the convenient upper bound with exponent scale

```
log X/log log X * (log log X/W_2(log X))^(45/4)
      *exp(-W_2(log X)/2).                          (G11)
```

Aggregation is first applied to the exact g in (G5); only afterward is sqrt(p) upper-bounded. No monotonicity of an arbitrary asymptotic replacement is needed, and w is never replaced inside exp(-w/2).

For a concrete classical instance, [MTY, Theorem 1.3] permits a=1/5.558691, and the density estimate recorded in [J] has A_d=5/2,B=1,C=0. For a VK instance, [MTY, Theorem 1.1] permits a=1/55.241, paired with the recorded Ivić density A_d=1600,B=3/2,C=15. Enlarge the fixed height cutoff to meet all hypotheses. These sufficient choices are not asserted to be optimal. Both specializations are unconditional consequences of the named ordinary-zeta theorems.

7. Attribution, comparison and references

The counting argument is a short adaptation of classical benefit and partition methods [N69, EN75, ENS]. Benefit inequalities, moving benefit allowances and prime counting near a cutoff have earlier precedents; they are not claimed as new here. The proof also uses classical integrated explicit-formula and density estimates. It assembles their stated inputs into the candidate-counting theorem above. The manuscript does not claim that this composition, or its final bound, is new.

The corrected Luca-Pomerance-Solé theorem [LPS] bounds the corresponding exception count beyond 7! by X^(O(1/log log X)). The finite convention change has no effect here. Relative to its exponent scale, the factor in our theorem is sqrt(P_0)exp(-W/2), which tends to zero because log P_0=o(W). This compares the displayed bounds only: it does not determine what sharper conclusions might already follow from their proof or other published results.

A completed bounded literature audit did not establish an exact prior statement or a demonstrated implication of the full theorem. This is not a novelty certificate. The Nicolas originals have now been recovered; their full implication analysis, and refined round-number counting behind LPS, remain open. The potential contribution is the spectral-budget derivation and its counting composition, subject to those comparisons. No effective numerical L or cutoff has been extracted. Unspecified constants do not establish inherent ineffectivity or vacuity at every concrete X. No lower bound for candidates, absence of exceptions, or RH conclusion follows from this upper bound.

- **[ENS]** P. Erdős, J.-L. Nicolas and A. Sárközy, *On Large Values of the Divisor Function*, Ramanujan Journal 2 (1998), pp.225-245; Lemmas 1-2 and the near-prime argument adjoining Theorem 6, especially pp.238-242. Publisher chapter record. Primary article scan in the proceedings volume. The partition theorem is imported; the full moving-budget integer theorem is not invoked.

- **[J]** D. R. Johnston, *Zero-density estimates and the optimality of the error term in the prime number theorem*, (2.2), discussion after (2.7) and (2.9). Version 1, arXiv:2411.13791v1. The log-free Jutila and Ivić density instances are imported as explicitly recorded in this version. Version 2 uses a different VK density witness; the present sufficient parameters are intentionally pinned to v1 and are not asserted optimal. Their original complete proofs were not re-audited. Johnston's pointwise PNT conclusion and its smoothness/unique-minimizer conditions are not inputs here.

- **[K]** A. Kumchev, *The Distribution of Prime Numbers*, Chapter 2, Exercise 16, printed p.38, and the Brun-Titchmarsh exercises. Author lecture notes. The primitive identity is stated there for psi_1(x)=sum_(n<=x)(x-n)Lambda(n). Since F=psi_1-x^2/2+2 and zeta'(0)/zeta(0)=log(2pi), it supplies the normalization in Section 3.2. It is a located classical import, not a proof of the explicit formula newly supplied here.

- **[MTY]** M. J. Mossinghoff, T. S. Trudgian and A. Yang, *Explicit zero-free regions for the Riemann zeta-function*, Theorems 1.1 and 1.3. Author paper.

- **[Z]** T. Zhao, *On the mean values of the error terms in Mertens' theorems*, Sections 1.1 and 2. Author paper. The Mertens normalization is recorded there. Its related tail expression is not substituted for the primitive identity, and its RH-dependent sign conclusions are not imported. The inspected results do not state our unconditional first-power W_2/P_0 estimate; that estimate is derived in Section 3 here. This attribution does not assert that no equivalent estimate exists elsewhere.

• **[LPS]** F. Luca, C. Pomerance and P. Solé, corrected Theorem 1 on Robin exceptions, *Corrigendum to #A8 of Volume 25, On Robin's Inequality*, Integers 25 (2025), Theorem 1. Corrected paper.

• **[Po]** C. Pomerance, *On the distribution of round numbers*, Lecture Notes in Mathematics 1122 (1985), pp.173-200, Theorem 6.1, p.195. Author copy. This is the slice-count estimate used by LPS. At their thresholds its stated upper bound alone gives exp(O(log X/log log X)), not the smaller scale required here. This observation does not exclude refined thresholds or stronger information about those slices.

• **[N69]** J.-L. Nicolas, *Ordre maximal d'un élément du groupe S_n des permutations et « highly composite numbers »*, Bulletin SMF 97 (1969), pp.129-191, Chapter 3, especially pp.154 and 167. Primary archive.

• **[EN75]** P. Erdős and J.-L. Nicolas, *Répartition des nombres superabondants*, Bulletin SMF 103 (1975), pp.65-90, Proposition 5, pp.73-74. Primary archive.

8. Formal verification and review scope

The mathematical proof in Sections 1-6 does not depend on a Lean implementation. The companion development verifies the finite encoding, arithmetic-benefit split, outside-coordinate bounds and distance budget. It also proves the finite partition inequality conditionally on explicitly stated Brun-Titchmarsh and ENS Lemma 1 hypotheses, with the correct one-based prime indexing.

Component	Status of the companion formal development
Injective support/multiplicity encoding and elementary outside/distance reduction	Proved in Lean; independently reviewed at recorded snapshots
Finite 16 Pi(T)^2 near-family bound and exp(F_out) composition	Proved in Lean conditional on BT and ENS Lemma 1; new implementation awaits independent review
BT and ENS Lemma 1 as unconditional Lean inputs	Not supplied by this route
ENS Lemma 2 asymptotic, F_out=o(G), and the spectral budget's (P1)	Not formalized for this route
Spectral tail and centered inclusion	Written deductions from the displayed inputs; not Lean-formalized
Global aggregation	A generic Lean theorem exists; the exact spectral specialization is not supplied

The fresh scoped build for the conditional finite interface used Lean 4.28.0, rebuilt 28 local modules and checked 105 printed theorem reports, each using only propext, Classical.choice and Quot.sound. The installed mathlib cache and toolchain remain trusted. This is not formal verification of Theorem 1 or of RH.

Codex assisted with proof development, implementation and manuscript maintenance; user-run Claude sessions provided adversarial model reviews. These are disclosed as model-assisted development and checking, not independent human peer review. The reports apply to their recorded source snapshots and do not certify this revised assembled manuscript. Specialist review of correctness and attribution is requested. The companion provenance and circulation-readiness notes record the revision and evidence boundaries.