The Redshift Bound from Quillen–Lichtenbaum

Shay Ben-Moshe

05/11/2026

In this talk I will explain a new proof of the following part of the redshift conjecture, which was originally proved by Clausen–Mathew–Naumann–Noel using the simultaneous purity theorem of Land–Mathew–Meier–Tamme.

Theorem 1. *For $\mathcal{C} \in \mathrm{Cat}_{\mathrm{st}}$ with L_n^f-local mapping spectra*

$$L_{\mathrm{T}(m)}\mathrm{K}(\mathcal{C}) = 0 \qquad \textit{for all } m \geq n{+}2.$$

1 Background

To start, let me remind you very briefly what is algebraic K-theory. Given $\mathcal{C} \in \mathrm{Cat}_{\mathrm{st}}$, the space of objects $\mathcal{C}^{\simeq}$ has the structure of a commutative monoid with respect to $\oplus$. Algebraic K-theory is obtained by splitting exact sequences, that is, informally

$$\mathrm{K}(\mathcal{C}) \;=\; \mathcal{C}^{\simeq} /\!/ \{[M] = [L] + [N] \;\mid\; L \to M \to N \text{ an exact sequence}\}\,.$$

Incidentally, this adds an inverse to every object by the sequence $M \to 0 \to \Sigma M$, so $\mathrm{K}(\mathcal{C})$ is a group-like commutative monoid, that is, a connective spectrum. The primary example is $\mathcal{C} = \mathrm{Perf}(R) = \mathrm{Mod}_R^{\omega}$ for a ring spectrum R, where we denote $\mathrm{K}(R) = \mathrm{K}(\mathrm{Perf}(R))$. Generally speaking, this is extremely difficult to compute, since we are adding plenty of relations. Chromatic homotopy theory, which I will now recall, offers a divide-and-conquer approach to spectra, which interacts in a surprising way with K-theory.

To start, let A be a (p-local) spectrum (or even just abelian group, working derived-ly). We can piece it from characteristic 0 and p information as the following pullback

$$\begin{array}{ccc} A & \longrightarrow & A_p \\ \downarrow & \lrcorner & \downarrow \\ A_{\mathbb{Q}} & \longrightarrow & (A_p)_{\mathbb{Q}} \end{array}$$

Chromatic homotopy theory refines this picture, further breaking down the p-completion. Very briefly, there is a filtration by a series of localizations

$$A \longrightarrow \cdots \longrightarrow L_n^f A \longrightarrow \cdots \longrightarrow L_1^f A \longrightarrow L_0^f A,$$

where $L_0^f A = A_{\mathbb{Q}}$, and L_1^f is the localization with respect to $\mathrm{KU}_{(p)}$. As above, we can also describe the relation between consecutive steps as some "completion", namely the localization with respect to some spectrum denoted $\mathrm{T}(n)$

$$\begin{array}{ccc} L_n^f A & \longrightarrow & L_{\mathrm{T}(n)} A \\ \downarrow & \lrcorner & \downarrow \\ L_{n-1}^f A & \longrightarrow & L_{n-1}^f L_{\mathrm{T}(n)} A \end{array}$$

From this perspective, $\mathrm{T}(n)$-local spectra are the building blocks of spectra. As in ordinary algebra, once we work locally at a prime, we can further simplify by passing to Galois extensions.

Theorem 2 (Burklund–Clausen–Levy)**.** *The algebraic closure of $\mathbb{S}_{\mathrm{T}(n)}$ is the Lubin–Tate spectrum, that is, it is the colimit over all finite Galois extensions*

$$\mathrm{E}_n = L_{\mathrm{T}(n)} (\operatorname*{colim}_{R/\mathbb{S}_{\mathrm{T}(n)}} R).$$

Remark 3. This (yet to be published) theorem builds on a series of results which may be more familiar to the experts in the $\mathrm{K}(n)$-local setting, due to Devinatz–Hopkins, Rognes, Baker–Richter, and Mathew. Using that instead, we get Theorem 1 for L_n-local categories.

The Lubin–Tate spectrum is indeed drastically simpler than the $\mathrm{T}(n)$-local sphere. For example, its homotopy groups are simply

$$\pi_* \mathrm{E}_n = W(\overline{\mathbb{F}}_p)[[u_1, \dots, u_{n-1}]][u^{\pm 1}], \qquad |u| = 2.$$

Thus, chromatic homotopy theory provides a way to decompose spectra. We can apply either to the input or the output of algebraic K-theory, and the redshift conjecture concerns the relation between the inside and outside localizations, loosely speaking saying that

algebraic K-theory increases chromatic height by 1.

One precise manifestation is Theorem 1, saying that for $\mathcal{C}$ with L_n^f-local mapping spectra, e.g. $\mathcal{C} = \mathrm{Perf}(R)$ for an L_n^f-local ring spectrum, the chromatic filtration on $\mathrm{K}(\mathcal{C})$ stops at L_{n+1}^f.

An important thing to mention is that the chromatic filtration on K-theory often doesn't converge. Nevertheless, sometimes the defect $\mathrm{fib}(\mathrm{K}(\mathcal{C})_{(p)} \to L_{n+1}^f \mathrm{K}(\mathcal{C}))$ is controllable. This is (closely related to) the Quillen–Lichtenbaum property mentioned in the title of the talk. A theorem of Hahn–Wilson (revisited by Angelini-Knoll) carries out a very explicit computation for a spectrum closely related to E_n, denoted $\mathrm{BP}\langle n \rangle$, showing that the fiber is bounded above. This gives an instance of Theorem 1.

Theorem 4 (Hahn–Wilson)**.** $\mathrm{K}_{\mathrm{T}(m)}(\mathrm{E}_n) = 0$ *for* $m \geq n+2$.

2 Proof of Theorem 1

The point of my paper is that Theorem 2 and Theorem 4 are sufficient to go all the way to general $\mathcal{C}$ with L_n^f-local mapping spectra, by the following steps

$$\mathrm{E}_n \quad \rightsquigarrow \quad R/\mathbb{S}_{\mathrm{T}(n)} \quad \rightsquigarrow \quad \mathbb{S}_{\mathrm{T}(n)} \quad \rightsquigarrow \quad L_n^f\mathbb{S} \quad \rightsquigarrow \quad \mathcal{C}.$$

In the proof, we will repeatedly use the silly observation that if M is a module over R, and $R = 0$, then $M = 0$. In particular if $R \to R'$ is a ring spectra map and $R = 0$ then $R' = 0$.

The proof proceeds by induction on the height n, starting from the case $n = -1$ where we take everything to be 0. The induction enters the proof through the second part of the following (which we in fact need twice).

Proposition 5. *K-theory sends the fracture square of a ring spectrum R to a pullback square*

$$\begin{array}{ccc} \mathrm{K}(L_n^f R) & \longrightarrow & \mathrm{K}(L_{\mathrm{T}(n)}R) \\ \downarrow & \lrcorner & \downarrow \\ \mathrm{K}(L_{n-1}^f R) & \longrightarrow & \mathrm{K}(L_{n-1}^f L_{\mathrm{T}(n)}R) \end{array}$$

Consequently, for $m \geq n+1$ the top map induces an isomorphism

$$\mathrm{K}_{\mathrm{T}(m)}(L_n^f R) \xrightarrow{\sim} \mathrm{K}_{\mathrm{T}(m)}(L_{\mathrm{T}(n)}R).$$

Proof sketch. The first holds since the vertical maps on Perf-s of the rings are Verdier quotients with the same kernels, and K-theory is localizing (alternatively, use Land–Tamme K-theory of pullbacks theorem and that L_{n-1}^f is smashing). The second part then follows from the inductive hypothesis. □

I learned the following argument from a paper of Elmanto–Nardin–Yang doing the $n = 0$ case (Mitchell's theorem).

Proposition 6. $\mathrm{K}_{\mathrm{T}(m)}(R) = 0$ *for some finite Galois extension $R/\mathbb{S}_{\mathrm{T}(n)}$.*

Proof (with two lies). Since K-theory commutes with filtered colimits, and $\mathrm{T}(m)$-localization commutes with all colimits, we get from Theorem 2 that*

$$\mathrm{K}_{\mathrm{T}(m)}(\mathrm{E}_n) = \operatorname*{colim}_{R/\mathbb{S}_{\mathrm{T}(n)}} \mathrm{K}_{\mathrm{T}(m)}(R).$$

All these K-theories are (commutative) ring spectra. Recall that a ring spectrum vanishes if and only if the elements 0 and 1 are the same. The left-hand side vanishes by Theorem 4. Since the sphere is compact*, $0 = 1$ happens at a finite stage of the colimit on the right-hand side, concluding the proof. □

Remark 7. The first problem is that E_n is the colimit over *R in $\mathrm{T}(n)$-local spectra*, not in spectra. Nonetheless, Proposition 5 (together with the fact that L_n^f-localization is smashing) implies that these two colimits have the same $\mathrm{T}(m)$-local K-theory. The second problem is that the outer colimit is *in $\mathrm{T}(m)$-local spectra*, not in spectra, and the $\mathrm{T}(m)$-local sphere is not compact. Nonetheless, we can work with $\mathrm{T}(m) \otimes -$ instead of $L_{\mathrm{T}(m)}$, and choose $\mathrm{T}(m)$ to be an $\mathbb{E}_1$-ring.

For the next step, we will use the following vanishing, developed with ChatGPT.

Proposition 8. *Let G be a finite group, let $\mathcal{C} \in \mathrm{CAlg}(\mathrm{Cat}_{\mathrm{perf}})^{\mathrm{B}G}$, and let Z be some spectrum.*

$$L_Z\mathrm{K}(\mathcal{C}) = 0 \qquad \Longleftrightarrow \qquad L_Z\mathrm{K}(\mathcal{C}^{hG}) = 0.$$

Proof idea. The $\Leftarrow$ part is easy, since $\mathcal{C}^{hG} \to \mathcal{C}$ is symmetric monoidal, inducing a ring map $L_Z\mathrm{K}(\mathcal{C}^{hG}) \to L_Z\mathrm{K}(\mathcal{C})$. For $\Rightarrow$ we use the multiplicative norm

$$\mathcal{C} \longrightarrow \mathcal{C}^{hG}, \qquad x \longmapsto \bigotimes_G gx,$$

which indeed lands in fixed points. This almost works: it is multiplicative but, despite sending 0 to 0 and 1 to 1, it is not additive, so it doesn't induce a map on K-theory. This can easily be remedied by instead considering G-tuples of elements (x_g) and sending them to $\bigotimes gx_g$, which is multi-linear, hence defines a map from $\bigotimes_G \mathcal{C}$. A theorem of Hilman–Ramzi / Cnossen–Haugseng–Lenz–Linskens shows that the norm is preserved by (Z-localized) K-theory, and upon taking (genuine) G-fixed points, we get a suitable map of ring spectra

$$(\bigotimes_G L_Z\mathrm{K}(\mathcal{C}))^G \longrightarrow L_Z\mathrm{K}(\mathcal{C}^{hG}). \qquad \square$$

Proposition 9. $\mathrm{K}_{\mathrm{T}(m)}(\mathbb{S}_{\mathrm{T}(n)}) = 0$.

Proof (with one lie). By Proposition 6, we have the vanishing for a finite Galois extension $R/\mathbb{S}_{\mathrm{T}(n)}$, so using Galois descent for modules[*], this follows from Proposition 8. $\square$

Remark 10. As before, the problem is that $R/\mathbb{S}_{\mathrm{T}(n)}$ is a Galois extension *in $\mathrm{T}(n)$-local spectra*, not in spectra, so $\mathrm{Perf}(R)^{hG} \neq \mathrm{Perf}(\mathbb{S}_{\mathrm{T}(n)})$. Rather, this holds for (dualizable) $\mathrm{T}(n)$-local modules. Luckily, the difference between (i.e. Verdier quotient of) perfect modules and dualizable $\mathrm{T}(n)$-local modules is L_{n-1}^f-local, so their $\mathrm{T}(m)$-local K-theory agrees by induction.

Proposition 11. $\mathrm{K}_{\mathrm{T}(m)}(L_n^f\mathbb{S}) = 0$.

Proof. This follows by combining Proposition 9 and the second part of Proposition 5. $\square$

Theorem 12. $\mathrm{K}_{\mathrm{T}(m)}(\mathcal{C}) = 0$ *for $\mathcal{C} \in \mathrm{Cat}_{\mathrm{st}}$ with L_n^f-local mapping spectra.*

Proof. Having L_n^f-local mapping spectra is equivalent to being a module over $\mathrm{Perf}(L_n^f\mathbb{S})$. Thus $\mathrm{K}_{\mathrm{T}(m)}(\mathcal{C})$ is a module over $\mathrm{K}_{\mathrm{T}(m)}(L_n^f\mathbb{S})$, and the latter vanishes by Proposition 11. $\square$