Descent and redshift in algebraic K-theory

Shay Ben-Moshe

03/06/2026

0 Zeroth K-group

Let R be a ring, and consider the collection of isomorphism classes of finitely generated projective R-modules (algebro-geometrically, these are vector bundles, but let us not dwell on this point). Direct sum makes this into a commutative monoid $(\mathrm{Proj}_R/\mathrm{iso}, \oplus) \in \mathrm{CMon}$.

Definition 0.1. We define the 0-th algebraic K group of R as the group-completion

$$\mathrm{K}_0(R) := (\mathrm{Proj}_R/\mathrm{iso}, \oplus)^{\mathrm{gpc}} \qquad \in \qquad \mathrm{Ab},$$

where the operation of group-completion is the left adjoint of the inclusion

$$\mathrm{Ab} \hookrightarrow \mathrm{CMon}.$$

Example 0.2. In a PID (e.g. fields and $\mathbb{Z}$) projective implies free, thus

$$\mathrm{K}_0(R) \simeq \{R^n \mid n \in \mathbb{N}\}^{\mathrm{gpc}} \simeq \mathbb{N}^{\mathrm{gpc}} \simeq \mathbb{Z}.$$

Example 0.3. Consider $R = \mathbb{Z}[\sqrt{-5}]$. It has a famous ideal $M = (2, 1 + \sqrt{-5})$, which is a projective module which isn't free. But one can quite directly show that $M \oplus M \simeq R^2$. In fact, this is all there is:

$$\mathrm{K}_0(\mathbb{Z}[\sqrt{-5}]) \simeq \mathbb{Z} \oplus \mathbb{Z}/2.$$

This is a reflection of the fact that $\mathbb{Z}[\sqrt{-5}]$ is not a unique factorization domain, as witnessed by $6 = 2 \cdot 3 = (1 + \sqrt{-5})(1 - \sqrt{-5})$.

Example 0.4. More generally, for a Dedekind domain we have

$$\mathrm{K}_0(R) \simeq \mathbb{Z} \oplus \mathrm{Pic}(R).$$

1 Algebraic K-theory

In our construction of $\mathrm{K}_0(R)$, we took the isomorphism classes of modules, neglecting the interesting information encoded in their automorphisms, i.e. in the groupoid $\mathrm{Proj}_R^{\simeq}$. I'll point out that for the purposes of the higher K-groups you may consider only the "boring" free modules R^n. While their structure modulo isomorphism is very simple (just $\mathbb{N}$, giving $\mathrm{K}_0 = \mathbb{Z}$), there are many interesting isomorphisms: indeed, the automorphisms of R^n is the group of invertible $n \times n$-matrices $\mathrm{GL}_n(R)$, which is an extremely interesting and rich object. Thus, we wish to retain all of this information. This is achieved by considering the groupoid $\mathrm{Proj}_R^{\simeq}$.

Importantly, of course, we can direct sum not only modules, but also maps between them – on matrices this corresponds to taking $A \in \mathrm{GL}_n(R)$ and $B \in \mathrm{GL}_k(R)$ we can form the block matrix $\begin{pmatrix} A & 0 \\ 0 & B \end{pmatrix} \in \mathrm{GL}_{n+k}(R)$. In other words, $\mathrm{Proj}_R^{\simeq}$ has a direct sum operation, making it into a commutative monoid in groupoids. We could have stopped there, but as I will explain in a moment, it turns out to be fruitful to consider it not only as a groupoid, but as an ∞-groupoid, i.e. a space or homotopy type

$$(\mathrm{Proj}_R^{\simeq}, \oplus) \qquad \in \qquad \mathrm{CMon}(\mathrm{Grpd}) \subset \mathrm{CMon}(\mathcal{S}).$$

We can then take as a definition the category of connective spectra to be commutative monoids in spaces, where each element has an inverse, which participates in an adjunction

$$\mathrm{CMon}(\mathcal{S}) \underset{\longleftarrow}{\xrightarrow{(-)^{\mathrm{gpc}}}} \mathrm{Ab}(\mathcal{S}) =: \mathrm{Sp}_{\geq 0}$$

Definition 1.1. We define the algebraic K-theory spectrum of R to be

$$\mathrm{K}(R) := (\mathrm{Proj}_R^{\simeq}, \oplus)^{\mathrm{gpc}} \qquad \in \qquad \mathrm{Sp}_{\geq 0},$$

and we define the K-groups as $\mathrm{K}_n(R) := \pi_n(\mathrm{K}(R))$.

A point that we will return to later is that this definition extends as-is to ring spectra.

One may wonder why should we care about ∞-groupoids, when $\mathrm{Proj}_R^{\simeq}$ is just a groupoid. The reason is that, just like a derived functor, the group completion may spread over all degrees. In general, K groups are notoriously difficult to calculate.

Example 1.2 (Bass–Milnor–Serre)**.** For a number field L we have $\mathrm{K}_1(\mathscr{O}_L) \simeq \mathscr{O}_L^{\times}$.

Essentially the only full computation is the following.

Example 1.3 (Quillen)**.** $\mathrm{K}_{2i}(\mathbb{F}_q) = 0$ and $\mathrm{K}_{2i-1}(\mathbb{F}_q) = \mathbb{Z}/(q^i - 1)$.

Example 1.4. We have

$$\mathrm{K}_1(\mathbb{Z}) = \mathbb{Z}^{\times} = \mathbb{Z}/2, \qquad \mathrm{K}_2(\mathbb{Z}) = \mathbb{Z}/2, \qquad \mathrm{K}_3(\mathbb{Z}) = \mathbb{Z}/48, \qquad \mathrm{K}_4(\mathbb{Z}) = 0.$$

The Vandiver conjecture is equivalent to $\mathrm{K}_{4n}(\mathbb{Z}) = 0$, and the second case $\mathrm{K}_8(\mathbb{Z}) = 0$ was only proven in 2018.

These exemplify the intricate arithmetic information encoded by K groups, for instance, they are closely related to special values of zeta functions. As another manifestation of this principle, we have the now-proven Quillen–Lichtenbaum conjecture, a consequence of Voevodsky–Rost's proof of the Bloch–Kato conjecture.

Theorem 1.5. *Let R be a suitable ring (regular Noetherian finitely generated) in which p is invertible. Then, there is a spectral sequence with*

$$E_2^{**} = H^*_{\text{et}}(R; \mathbb{Z}/p(\frac{*}{2}))$$

which converges to $\pi_*(\mathrm{K}(R)/p)$ for $* \gg 0$.

Loosely speaking, there is some relation between K-theory and etale cohomology. This statement, however, is somewhat odd, with the connection holding only for $* \gg 0$. This leads us to our next topic.

2 Descent

As we have seen, algebraic K-theory is a fairly intricate invariant, and is difficult to compute. One way to approach such calculations is via descent.

Theorem 2.1 (Thomason)**.** *K-theory satisfies Zariski descent. That is, for affine schemes $X = U \cup V$, we have a pullback square in spectra*

$$\begin{array}{ccc} \mathrm{K}(X) & \longrightarrow & \mathrm{K}(U) \\ \downarrow & \lrcorner & \downarrow \\ \mathrm{K}(V) & \longrightarrow & \mathrm{K}(U \cap V) \end{array}$$

Remark 2.2. Note that the pullback must be taken in spectra (like a derived pullback), as with taking mod p above. It is *not* true that each K_n satisfies descent. In more down to earth terms though, this gives rise to a long exact sequence on K groups.

In fact Thomason proved more than that: K-theory satisfies Nisnevich descent. The Nisnevich topology is intermediate between Zariski and etale. A map $\coprod U_i \to X$ is a Nisnevich cover if it is an etale cover and for each point $x \in X$ there exists some i and $u \in U_i$ s.t. $k(x) \xrightarrow{\sim} k(u)$ is an isomorphism. Notably, K-theory does *not* satisfy etale descent, precisely because it fails to satisfy Galois descent. That is, there are examples of a G-Galois extension $L \to L'$ such that

$$\mathrm{K}(L) \longrightarrow \mathrm{K}(L')^{hG}$$

is not an isomorphism. This ties back to the peculiarities of the Quillen–Lichtenbaum conjecture, which can be remedied using chromatic homotopy theory.

3 Chromatic localizations

A very useful paradigm in ordinary algebra is studying questions one prime at a time and then gluing the results. For simplicity, let us work p-locally, then this decomposition is controlled by the fairly simple topological space

$$\mathrm{Spec}(\mathbb{Z}_{(p)}) = \{(0) \to (p)\}.$$

For example, a p-local abelian group A can be recovered from its rationalization and p-completion glued along the rationalization of the p-completion

$$A = A_{\mathbb{Q}} \underset{(A_p)_{\mathbb{Q}}}{\times} A_p.$$

Surprisingly, this picture refines in spectra – this is the subject of chromatic homotopy theory. The p-completion further decomposes into infinitely many "new characteristics"

$$\mathrm{Spec}(\mathbb{S}_{(p)}) = \{(0) \to (p,1) \to \cdots \to (p,n) \to \cdots \to (p,\infty)\}.$$

In particular, there is a localization $L_{\mathrm{T}(n)} \colon \mathrm{Sp} \to \mathrm{Sp}_{\mathrm{T}(n)}$ for every *height* $n \geq 0$ (and implicit prime p), where the case of height 0 is rationalization, while the higher heights are refinements of the p-completion (we remark that there is also some information not captured by any of these heights, but for our purposes it will be negligible).

I will not give too many details, but let me comment that this comes from a very close connection between spectra and (1-dimensional) formal groups, which over $\overline{\mathbb{F}}_p$ are classified by a number called the height, due to deep insights of Quillen, Morava, Ravenel and Devinatz–Hopkins–Smith, among many others.

Example 3.1. This is a purely homotopical phenomenon, in the sense that given an abelian group A, considered as a spectrum, we have $L_{\mathrm{T}(n)}A = 0$ for $n \geq 1$.

Example 3.2. Tautologically, the unit is the $\mathrm{T}(n)$-local sphere $L_{\mathrm{T}(n)}\mathbb{S}$.

Example 3.3. Topological K-theory KU is of height $n = 1$, in the sense that its $\mathrm{T}(1)$-localization is already the entire p-completion $L_{\mathrm{T}(1)}\mathrm{KU} = \mathrm{KU}_p$. Importantly, this is the Galois closure of $L_{\mathrm{T}(1)}\mathbb{S}$.[1]

Example 3.4. The Lubin–Tate deformation theory of formal groups can be carried out in spectra: given a formal group $\mathbf{G}$ of height n over a field L, there is a commutative ring spectrum $\mathrm{E}(L, \mathbf{G})$, whose π_0 is the ordinary Lubin–Tate ring, which is $\mathrm{T}(n)$-local $\mathrm{E}(L, \mathbf{G}) = L_{\mathrm{T}(n)}\mathrm{E}(L, \mathbf{G})$. For $n = 1$, taking the multiplicative formal group $\mathbf{G} = \mathbf{G}_m$ recovers the previous example $\mathrm{E}(\mathbb{F}_p, \mathbf{G}_m) = \mathrm{KU}_p$. Importantly, this is the Galois closure of $L_{\mathrm{T}(n)}\mathbb{S}$ (due to Devintaz–Hopkins, Baker–Richter and Burklund–Clausen–Levy).[2]

Going back to algebraic K-theory, we have the following fundamental theorem.

Theorem 3.5 (Mitchell)**.** *For any ring R and $n \geq 2$ we have $L_{\mathrm{T}(n)}\mathrm{K}(R) = 0$.*

Contrast this with the case of ordinary abelian group: whereas there $L_{\mathrm{T}(n)}A = 0$ for $n \geq 1$, here we see information at heights 0 and 1. Furthermore, Thomason showed that the $\mathrm{T}(1)$-local part is precisely the etale part.

Theorem 3.6 (Thomason)**.** *The functor $L_{\mathrm{T}(1)}\mathrm{K}(-)$ satisfies Galois, hence etale, descent. In fact, it is the etale sheafification for p-invertible rings, for which we have a spectral sequences*

$$E_2^{**} = H^*_{\text{et}}(R; \mathbb{Z}_p(\frac{*}{2})) \implies \pi_*(L_{\mathrm{T}(1)}\mathrm{K}(R)).$$

[1] Up to adding roots of unity.

[2] When we take $L = \overline{\mathbb{F}}_p$.

Thus the Quillen–Lichtenbaum conjecture can be reformulated as saying that

$$\pi_*(\mathrm{K}(R)) \longrightarrow \pi_*(L_{\mathrm{T}(1)}\mathrm{K}(R))$$

is an isomorphism for $* \gg 0$. This demonstrates the deep connection between algebraic K-theory and chromatic homotopy.

4 Redshift

We have seen various phenomena in the case of ordinary rings, namely rings of height 0 (as $L_{\mathrm{T}(n)}R = 0$ for $n \geq 1$). Based on this case, as well as some computational evidence for ring spectra of height 1, in the early 2000's Ausoni–Rognes proposed a cluster of conjectures, now known as the redshift conjecture, which have seen tremendous breakthroughs in recent years by works of various groups, including Hahn–Wilson, Land–Mathew–Meier–Tamme, Clausen–Mathew–Naumann–Noel and Burklund–Schlank–Yuan.

First, and perhaps most notably, we saw that the K-theory of ordinary rings (height 0) has height 1. A generalization of this is what gives redshift its name: algebraic K-theory increases height by 1.

Theorem 4.1. *K-theory increases chromatic height by one: let R be a $\mathrm{T}(n)$-local commutative ring spectrum then*

$$L_{\mathrm{T}(n+1)}\mathrm{K}(R) \neq 0, \qquad L_{\mathrm{T}(m)}\mathrm{K}(R) = 0 \quad \forall m \geq n+2.$$

Second, we saw that the localization at height 1 has better descent properties.

Theorem 4.2. *$\mathrm{T}(n+1)$-localized K-theory satisfies Galois descent for finite p-groups. That is, let $R \to S$ be a Galois extension of $\mathrm{T}(n)$-local rings with finite p-group Galois group G, then*

$$L_{\mathrm{T}(n+1)}\mathrm{K}(R) \xrightarrow{\sim} L_{\mathrm{T}(n+1)}\mathrm{K}(S)^{hG}.$$

Finally, there are higher chromatic analogues of Quillen–Lichtenbaum, known only in examples (which would require too much setup for the time frame of this talk).

5 Cyclotomic redshift and the telescope conjecture

I'd like to discuss joint work with Carmeli, Schlank and Yanovski. Our main theorem is an extension of this last result in a somewhat abstract direction, replacing the Galois group by a π-finite group (group in spaces such that the entire π_* is finite). Instead of describing it, I will formulate some consequences that are somewhat more tangible, concerning analogues of cyclotomic extension.

Given a commutative ring R, we may consider the cyclotomic extension $R[\omega_{p^k}]$. We note that this construction works well away from p, i.e. when p is invertible in R, for instance it is a $(\mathbb{Z}/p^k)^\times$-Galois extension, while for $\mathbb{F}_p$-algebras it is totally ramified, and behaves completely differently. Surprisingly, while $\mathrm{T}(n)$-local spectra are in a sense at the prime p, they still have a working theory of cyclotomic extensions, due to Carmeli–Schlank–Yanovski (this uses higher semiadditivity, where the aforementioned π-finite groups satisfy Poincare duality as though they were 0-dimensional oriented manifolds). That is, for a $\mathrm{T}(n)$-local commutative ring spectrum R, there is a $(\mathbb{Z}/p^k)^\times$-Galois extension $R[\omega_{p^r}^{(n)}]$, called the height n cyclotomic extension. One of our main results is the following.

Theorem 5.1 (B.-M.–Carmeli–Schlank–Yanovski)**.** *$\mathrm{T}(n+1)$-localized K-theory sends height n cyclotomic extension to height $n+1$ cyclotomic extensions, i.e.*

$$L_{\mathrm{T}(n+1)}\mathrm{K}(R[\omega_{p^r}^{(n)}]) \simeq L_{\mathrm{T}(n+1)}\mathrm{K}(R)[\omega_{p^r}^{(n+1)}]$$

together with the $(\mathbb{Z}/p^k)^\times$-action.

Remark 5.2. Note that $(\mathbb{Z}/p^k)^\times$ is not a p-group (e.g. $|(\mathbb{Z}/p)^\times| = p-1$), and hence doesn't fit in the previous theorem.

Finally, I would like to discuss Galois descent for profinite groups. Given a G-Galois extension S of R, we have $R \xrightarrow{\sim} S^{hG}$. In ordinary algebra this holds not only for finite Galois extensions but also for profinite extensions (system of finite Galois extensions). In higher algebra this may a priori *fail*: S may not be faithful over R (while each of the finite Galois extensions is). Since passing to a Galois extension simplifies things a lot, it is very helpful to know whether this holds or not. For example, while we know almost nothing about the unit $L_{\mathrm{T}(n)}\mathbb{S}$, we know a lot about its algebraic closure $\mathrm{E}_n = \mathrm{E}(\overline{\mathbb{F}}_p, \mathbf{G})$, such as its homotopy groups, and at height $n \leq 2$ we even completely understand the action of the (profinite) Galois group. While originally stated in a different language, this is the subject of Ravenel's long-standing telescope conjecture:

Conjecture 5.3. *E_n is faithful over $L_{\mathrm{T}(n)}\mathbb{S}$.*

This was known at heights 0 (trivial) and 1 (Mahowald, Miller), but was otherwise open for 40 years.

Our theorem leads to an approach for *disproving* the conjecture at height $n+1$, namely showing that E_{n+1} is not faithful over $L_{\mathrm{T}(n+1)}\mathbb{S}$. Since E_{n+1} is the algebraic closure, it demonstrate that for some intermediate extension. In particular, it suffices to find some ring R of height n, such that $L_{\mathrm{T}(n+1)}\mathrm{K}(R)[\omega_{p^\infty}^{(n+1)}]$ does not descend to $L_{\mathrm{T}(n+1)}\mathrm{K}(R)$. Our theorem identifies it with $L_{\mathrm{T}(n+1)}\mathrm{K}(R[\omega_{p^\infty}^{(n)}])$, importantly, $\mathbb{Z}_p^\times$-equivariantly – translating the problem to understanding the action on the height $n+1$ algebraic K-theory of something coming from height n. In their remarkable paper, Burklund–Hahn–Levy–Schlank disproved the telescope conjecture by finding a suitable R and for which hyperdescent along the cyclotomic tower fails.

Theorem 5.4. *The telescope conjecture is false for $n \geq 2$.*

6 Loose ends and further directions

6.1 Higher semiadditivity

A key construction that appeared above is the higher height analogues of cyclotomic extensions. As I have briefly mentioned, their construction uses a remarkable phenomenon called higher semiadditivity. π-finite spaces (namely truncated and each π_i is finite), such as the classifying space of a finite group $\mathrm{B}G$, are generally very complicated homologically. For example, they have infinitely many cells, and their homology is non-vanishing in infinitely many degrees. Nevertheless, their $\mathrm{T}(n)$-homology is much more controlled: they satisfy Poincare duality, behaving as though they were 0-dimensional oriented manifolds.

Theorem 6.1 (Hopkins–Lurie, Carmeli–Schlank–Yanovski, B.-M.)**.** *Let R be a $\mathrm{T}(n)$-local spectrum and let X be a π-finite space, then there is a canonical isomorphism between the $\mathrm{T}(n)$-local homology and cohomology*

$$R[X] \simeq R^X \qquad \in \qquad \mathrm{Sp}_{\mathrm{T}(n)}.$$

This was first proved by Hopkins–Lurie, for R satisfying descent over E_n, by a very direct computation over E_n. This was then proved for all R by Carmeli–Schlank–Yanovski, by a much "softer" argument. This result in this latter level of generality was crucial for the arguments I have outlined in the previous section. More recently, I reversed the logic, and used the K-theoretic techniques similar to the above to give a new proof of the result, under some intermediary assumptions, for example recovering the Hopkins–Lurie result but also strong enough for the application to the telescope conjecture. Very briefly put, the idea is to show that $\mathrm{T}(n+1)$-localized K-theory takes the self-duality of homology at height n to a self-duality data at height $n+1$, giving a proof by induction on the height.

It should be noted that while the theorem holds for every n separately, it is *not* true if R is supported on more than one height (that is, not $\mathrm{T}(n)$-local for any n). Nevertheless a connection between higher semiadditivity across different chromatic heights can be established, as was proven by Lurie's tempered ambidexterity theorem and my follow up work on the subject.

6.2 Trace methods (and higher semiadditivity)

One of the key approaches for computing or getting a handle on algebraic K-theory, which I left implicit in the discussion so far, goes under the name of trace methods. Very briefly put, the idea is to study $\mathrm{K}(R)$, that is, modules and morphisms between them, by their dimensions and traces. This is a collection of invariants approximating $\mathrm{K}(R)$, such as topological cyclic homology and topological Hochschild homology

$$\mathrm{K}(R) \longrightarrow \mathrm{TC}(R) \longrightarrow \mathrm{THH}(R).$$

The main advantage of these invariants is that they are much easier to compute: they are built from R itself, rather than its category of modules, by a series of approachable operations such as tensor products and (co)limits. Moreover, they offer an excellent approximation, in fact, $\mathrm{K}(R) \to \mathrm{TC}(R)$ is $\mathrm{T}(m)$-local isomorphism for any $m \geq 2$!

In light of the relation between algebraic K-theory and higher semiadditivity, it is natural to try and study the behavior of higher semiadditivity under these invariants. They have a remarkable similarity to the connection I've mentioned across different chromatic heights, and I am presently working on relating the two more directly.