DUAL MODULES

KEITH CONRAD

1. Introduction

Let R be a commutative ring. For two (left) R-modules M and N, the set $\mathrm{Hom}_R(M, N)$ of all R-linear maps from M to N is an R-module under natural addition and scaling operations on linear maps. (If R were non-commutative then the definition $(r \cdot f)(m) = r \cdot (f(m))$ would yield a function $r \cdot f$ from M to N which is usually not R-linear. Try it!) In the special case where $N = R$ we get the R-module

$$M^\vee = \mathrm{Hom}_R(M, R).$$

This is called the *dual module*, *dual space*, or R-dual of M. Elements of $M^\vee$ are called *linear functionals* or simply *functionals*.

Here are some places in mathematics where dual modules show up:

(1) linear algebra: coordinate functions on $\mathbf{R}^n$.
(2) analysis: integration on a space of continuous functions.
(3) geometry: the abstract definition of a tangent space (directional derivatives).
(4) number theory: the different ideal of a number field.

There really is no picture of the dual module, but its elements could be thought of as "potential coordinate functions" on M (plus the function 0, so we have a module). This idea is accurate if M if a finite-dimensional vector space, or even a finite-free module, but in more general settings it can be an oversimplification.

In Section 2 we will look at some examples of dual modules. The behavior of the dual module on direct sums and direct products is the topic of Section 3. The special case of dual modules for finite free modules is in Section 4, where we meet the important double duality isomorphism. Section 5 describes the construction of the dual of a linear map between modules, which generalizes the matrix transpose. In Section 6 we will see how dual modules arise in concrete ways using the language of (perfect) pairings.

2. Examples

Example 2.1. What are the functionals on $\mathbf{R}^n$? Examples are the standard coordinate functions on $\mathbf{R}^n$:

$$c_1e_1 + \cdots + c_ne_n \mapsto c_i.$$

More generally, dotting on $\mathbf{R}^n$ with a fixed vector is in the dual space: for each $v \in \mathbf{R}^n$, let $\varphi_v \colon \mathbf{R}^n \to \mathbf{R}$ by

$$\varphi_v(w) = v \cdot w.$$

The standard coordinate functions on $\mathbf{R}^n$ arise this way when v is one of the standard basis vectors $e_1, \ldots, e_n$ of $\mathbf{R}^n$. We will show $\mathbf{R}^n \cong (\mathbf{R}^n)^\vee$ by $v \mapsto \varphi_v$. That is, the functionals on $\mathbf{R}^n$ are dot products with different vectors in $\mathbf{R}^n$.

Each φ_v is linear, so it lies in $(\mathbf{R}^n)^\vee$. Moreover, since $\varphi_{v+v'} = \varphi_v + \varphi_{v'}$ and $\varphi_{cv} = c\varphi_v$ for $c \in \mathbf{R}$ (just check both sides dot any $w \in \mathbf{R}^n$ in the same way), sending v to φ_v is a linear

map $\mathbf{R}^n \to (\mathbf{R}^n)^\vee$. This is injective since if $\varphi_v = 0$ in $(\mathbf{R}^n)^\vee$ then $v \cdot w = 0$ for all $w \in \mathbf{R}^n$, and taking $w = e_1, \ldots, e_n$ shows $v = 0$. To show surjectivity, pick an $f \in (\mathbf{R}^n)^\vee$. Then, for any $w = (c_1, \ldots, c_n) = \sum c_i e_i$ in $\mathbf{R}^n$,

$$\begin{aligned} f(w) &= f(\sum c_i e_i) \\ &= \sum c_i f(e_i) \\ &= (c_1, \ldots, c_n) \cdot (f(e_1), \ldots, f(e_n)) \\ &= \varphi_v(w) \end{aligned}$$

where $v = (f(e_1), \ldots, f(e_n))$. So $f = \varphi_v$ for this choice of v.

The fact that $\mathbf{R}^n$ can be identified with $(\mathbf{R}^n)^\vee$ using the dot product may have delayed somewhat the development of abstract linear algebra, since it takes a certain amount of insight to realize that the dual space is an object of independent interest when it is nothing really new in the classical setting of Euclidean space. That dual spaces are something separate from the original space was first recognized in functional analysis, where for instance the dual space of a space of continuous functions is a space of measures.

Example 2.2. For any R, consider R^n ($n \geq 1$) as an R-module in the usual way. The dot product maps $R^n \times R^n$ to R, every element of $(R^n)^\vee$ has the form $\varphi_v(w) = v \cdot w$ for a unique $v \in R^n$, and the correspondence $v \mapsto \varphi_v$ is an R-module isomorphism from R^n to $(R^n)^\vee$. The proof is just like the case of $\mathbf{R}^n$.

In particular, $R^\vee = \mathrm{Hom}_R(R, R)$ is isomorphic to R in the sense that every R-linear map $R \to R$ has the form $\varphi_a(r) = ar$ for a unique $a \in R$. The isomorphism $R^\vee \cong R$ is $\varphi \mapsto \varphi(1)$ and, in the other direction, $a \mapsto \varphi_a$.

Example 2.3. If $M = 0$ then $M^\vee = 0$ too (the only linear map $0 \to R$ is the zero map).

Theorem 2.4. *If M is a finite free R-module of rank n, i.e., $M \cong R^n$ as R-modules, then $M^\vee$ is finite free of rank n too.*

Proof. The case $n = 0$ (so $M = 0$) is trivial, so we may assume $n \geq 1$. Since M is isomorphic to R^n, the dual module $M^\vee$ is isomorphic to $(R^n)^\vee$. Explicitly, letting $L\colon M \to R^n$ be an R-module isomorphism, $(R^n)^\vee \cong M^\vee$ as R-modules by $\varphi \mapsto \varphi \circ L$. Here is a picture:

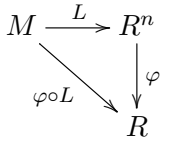

By Example 2.2, $(R^n)^\vee \cong R^n$, so $M^\vee \cong R^n$. □

The proof of Theorem 2.4 depended on the choice of an R-basis of M when we introduced an isomorphism $M \cong R^n$. If we change the basis then the isomorphism of $M^\vee$ with R^n changes. So it is a coordinate-free fact that $M^\vee \cong M$ as R-modules when M is finite free, but there isn't a *natural* isomorphism of $M^\vee$ with M in general.

Example 2.5. Let $R = \mathbf{Z}$, so R-modules are abelian groups. For an abelian group A, its $\mathbf{Z}$-dual is $A^\vee = \mathrm{Hom}_{\mathbf{Z}}(A, \mathbf{Z})$. If $A = \mathbf{Z}^n$, we can identify $A^\vee$ with A using dot products with varying n-tuples in A, just as over the reals in Example 2.1. On the other hand, if we take $A = \mathbf{Q}$ and treat it as a $\mathbf{Z}$-module (not as a $\mathbf{Q}$-vector space!) then $\mathbf{Q}^\vee = \mathrm{Hom}_{\mathbf{Z}}(\mathbf{Q}, \mathbf{Z})$ is zero: when $f \in \mathbf{Q}^\vee$, for any $r \in \mathbf{Q}$ the *integer* $f(r)$ satisfies $f(r) = 2^n f(r/2^n)$ with

$f(r/2^n) \in \mathbf{Z}$, so $f(r)$ is divisible by arbitrarily high powers of 2. Thus $f(r) = 0$ for all r, so $f = 0$. If we treat $\mathbf{Q}$ as a $\mathbf{Q}$-vector space then $\mathbf{Q}^\vee = \mathrm{Hom}_{\mathbf{Q}}(\mathbf{Q}, \mathbf{Q})$ is not zero (it is isomorphic to $\mathbf{Q}$). The notation $M^\vee$ for the dual module leaves out reference to the ring R over which M is an R-module. We could write, say, $M^{\vee_R}$ if we want to put R in the notation, but we will generally just rely on context.

If A is a finite abelian group, its $\mathbf{Z}$-dual is 0 since a group homomorphism takes elements of finite order to elements of finite order and the only element of finite order in $\mathbf{Z}$ is 0.

Remark 2.6. Since the $\mathbf{Z}$-dual of a finite abelian group A is zero, and thus uninteresting, there is an alternate notion of dual for finite abelian groups called the (Pontryagin) *dual group*: $A^\wedge = \mathrm{Hom}(A, S^1)$, the set of group homomorphisms from A to the unit circle $S^1 \subset \mathbf{C}^\times$ under pointwise multiplication. Another common notation for this dual group is $\widehat{A}$. The dual group is used in the study of characters on finite abelian groups. Generalizing finite abelian groups to locally compact abelian groups, the dual group becomes a central object of Fourier analysis on groups.

Theorem 2.7. *Let R be an integral domain with fraction field K, and let M be a nonzero R-module in K. Then $M^\vee = \mathrm{Hom}_R(M, R)$ is isomorphic to $\{c \in K : cM \subset R\}$.*

An example of M to keep in mind is a nonzero ideal of R.

Proof. For any $c \in K$ such that $cM \subset R$, the function $\varphi_c \colon x \mapsto cx$ is an R-linear map from M to R. Conversely, let $\varphi \colon M \to R$ be R-linear. We will construct a $c \in K$ such that $\varphi(x) = cx$ for all $x \in M$. It will then follow that $cM = \varphi(M) \subset R$, so every element of $M^\vee$ arises by our concrete construction.

Fix a nonzero $m_0 \in M$. For $x \in M$, write m_0 and x as ratios in R with a common denominator: $m_0 = a/d$ and $x = b/d$, where $a, b, d \in M$. Since φ is R-linear,

$$d\varphi(m_0 x) = \varphi(dm_0 x) = \varphi(ax) = a\varphi(x) = dm_0\varphi(x)$$

and

$$d\varphi(m_0 x) = \varphi(dm_0 x) = \varphi(bm_0) = b\varphi(m_0) = d\varphi(m_0)x,$$

so $\varphi(x) = cx$, where $c = \varphi(m_0)/m_0 \in K$.

Let $\{c \in K : cM \subset R\} \to M^\vee$ by $c \mapsto \varphi_c$. Since $\varphi_{c+c'} = \varphi_c + \varphi_{c'}$ and $\varphi_{rc} = r\varphi_c$ for $r \in R$, our mapping is R-linear. We showed above that it is surjective, and since we are working a field the mapping is injective, and thus $M^\vee \cong \{c \in K : cM \subset R\}$. □

Example 2.8. Let $R = \mathbf{Z}[\sqrt{-14}]$ and I be the ideal $(3, 1 + \sqrt{-14})$ in R. By Theorem 2.7, the elements of $I^\vee = \mathrm{Hom}_R(I, R)$ can be regarded as $\{c \in \mathbf{Q}(\sqrt{-14}) : cI \subset R\}$. We will describe this set in terms of the ideal $J = (3, 1 - \sqrt{-14})$. Verify as an exercise that

$$IJ = (3) = 3R. \tag{2.1}$$

Therefore if $x \in I$ and $y \in J$, then $xy \in 3R$, so $xy/3 \in R$. Conversely, if $y \in \mathbf{Q}(\sqrt{-14})$ satisfies $yI \subset R$, then $yIJ \subset RJ = J$, so $3yR \subset J$, which implies $y \in (1/3)J$. Thus $I^\vee \cong (1/3)J$ where $t \in (1/3)J$ acts on I by $x \mapsto tx$.

We can also view $I^\vee$ as J by associating to each $y \in J$ the linear map $I \to R$ where $x \mapsto xy/3$.

By a similar argument $J \cong I^\vee$ as R-modules (every element of $J^\vee$ has the form $y \mapsto xy/3$ for a unique $x \in I$), so I and J can each be viewed as the R-dual of the other.

Example 2.9. When $R = \mathbf{Z}[X]$ and I is the maximal ideal $(2, X)$,

$$I^\vee = \{f \in \mathbf{Q}(X) : fI \subset R\} = \{f \in \mathbf{Q}(X) : 2f \in \mathbf{Z}[X] \text{ and } Xf \in \mathbf{Z}[X]\} = \mathbf{Z}[X] = R$$

since $\mathbf{Z}[X]$ is a UFD and 2 and X are relatively prime. That is, the only $\mathbf{Z}[X]$-linear maps $(2, X) \to \mathbf{Z}[X]$ are the multiplication maps $g \mapsto fg$ for $f \in \mathbf{Z}[X]$.

Corollary 2.10. *Let R be an integral domain with fraction field K. A nonzero R-module M in K has nonzero dual module if and only if M admits a common denominator from R: there is some $d \in R - \{0\}$ such that $M \subset (1/d)R$.*

Proof. If $M \subset (1/d)R$ for some $d \in R - \{0\}$ then $dM \subset R$, so $x \mapsto dx$ is an example of a nonzero element of $M^\vee$.

Conversely, assume $M^\vee$ is nonzero. Then there is some $c \in K^\times$ such that $cM \subset R$. Write $c = a/b$ where a and b are in $R - \{0\}$. Then $aM \subset bR \subset R$, so a is a common denominator for M. □

This "explains" why we found in Example 2.5 that the $\mathbf{Z}$-dual of $\mathbf{Q}$ is zero: $\mathbf{Q}$ as a $\mathbf{Z}$-module does not admit a common denominator from $\mathbf{Z}$.

Corollary 2.10 tells us that it is natural to look at the R-modules in K that admit a common denominator; for other R-modules in K, the R-dual is 0.

Theorem 2.11. *Let R be an integral domain with fraction field K. For two nonzero R-modules M and N in K, the R-module $\mathrm{Hom}_R(M, N)$ is isomorphic to $\{c \in K : cM \subset N\}$. If M and N admit a common denominator from R then so does $\mathrm{Hom}_R(M, N)$. Equivalently, if $M^\vee$ and $N^\vee$ are nonzero then $\mathrm{Hom}_R(M, N) \neq 0$.*

Proof. Exercise. □

3. Duals, Direct Sums, and Direct Products

The dual module construction behaves nicely on direct sums.

Theorem 3.1. *For R-modules M and N, $(M \oplus N)^\vee \cong M^\vee \oplus N^\vee$.*

Proof. Given $\varphi \in (M \oplus N)^\vee$, we obtain elements $f \in M^\vee$ and $g \in N^\vee$:

$$f(m) = \varphi(m, 0), \quad g(n) = \varphi(0, n).$$

Sending φ to the ordered pair (f, g) is an R-linear map from $(M \oplus N)^\vee$ to $M^\vee \oplus N^\vee$. Conversely, given $(f, g) \in M^\vee \oplus N^\vee$, set $\varphi\colon M \oplus N \to R$ by $\varphi(m, n) = f(m) + g(n)$. Then $\varphi \in (M \oplus N)^\vee$ and we have constructed an R-linear map $M^\vee \oplus N^\vee \to (M \oplus N)^\vee$ which is inverse to our map in the other direction. □

Example 3.2. Let $A = \mathbf{Z} \oplus (\mathbf{Z}/2\mathbf{Z})$. Viewing A as a $\mathbf{Z}$-module in the obvious way, $A^\vee \cong \mathbf{Z}^\vee \oplus (\mathbf{Z}/2\mathbf{Z})^\vee = \mathbf{Z}^\vee$ by Example 2.5. Using Example 2.2, which says $\mathbf{Z}^\vee \cong \mathbf{Z}$ using multiplication maps, $A^\vee$ consists of the functions $f_k(x, y) = kx$ for different integers k.

Since the direct sum of modules is associative (up to isomorphism), Theorem 3.1 extends by induction to finite direct sums of any length: the dual of any finite direct sum of modules is naturally isomorphic to the direct sum of the dual modules.

What can we say about the dual of a direct sum of infinitely many R-modules? It is *not* isomorphic to the direct sum of the dual modules. It's isomorphic to their direct product. Pay attention in the following proof to the different meanings of direct sums and direct products.

Theorem 3.3. *Let M_i $(i \in I)$ be a family of R-modules. There is an isomorphism $(\bigoplus_{i\in I} M_i)^\vee \cong \prod_{i\in I} M_i^\vee$*

Proof. Let $\varphi \in (\bigoplus_{i\in I} M_i)^\vee$, so $\varphi\colon \bigoplus_{i\in I} M_i \to R$ is R-linear. Viewing M_i as a submodule of $\bigoplus_{i\in I} M_i$ in the usual way, the restriction $\varphi|_{M_i}$ is an R-linear map from M_i to R. The entire collection of restrictions $(\varphi|_{M_i})_{i\in I}$ lies in the direct product $\prod_{i\in I} M_i^\vee$. There is no reason to expect most restrictions $\varphi|_{M_i}$ are identically 0, so the collection of restrictions is usually *not* in the direct sum of the $M_i^\vee$'s.

We have a map $(\bigoplus_{i\in I} M_i)^\vee \to \prod_{i\in I} M_i^\vee$ given by $\varphi \mapsto (\varphi|_{M_i})_{i\in I}$. This is R-linear (check!). We will write down the inverse map. Given a $(\psi_i)_{i\in I} \in \prod_{i\in I} M_i^\vee$, define $\psi \in (\bigoplus_{i\in I} M_i)^\vee$ by

$$\psi((m_i)_{i\in I}) = \sum_{i\in I} \psi_i(m_i).$$

That is, ψ of an element $(m_i) \in \bigoplus_{i\in I} M_i$ is the sum of each ψ_i at the i-th coordinate m_i of the element. This sum is a finite sum because all but finitely many m_i's are *zero*, so $\psi_i(m_i) = 0$ for all but finitely many $i \in I$ (even if ψ_i is not the function 0). The reader can check ψ is R-linear, so it is in the dual module of $\bigoplus_{i\in I} M_i$. Sending $(\psi_i)_{i\in I}$ to ψ is an R-linear map from $\prod_{i\in I} M_i^\vee$ to $(\bigoplus_{i\in I} M_i)^\vee$. It is left to the reader to check this is a 2-sided inverse to the map we constructed in the other direction. □

Remark 3.4. It was irrelevant in the above proof that we were working with dual modules (values in R). The same proof shows, for any R-module N and family of R-modules M_i that $\mathrm{Hom}_R(\bigoplus_{i\in I} M_i, N) \cong \prod_{i\in I} \mathrm{Hom}_R(M_i, N)$. Taking $N = R$ recovers Theorem 3.3.

Example 3.5. Let $R = \mathbf{Z}$ and let $M = \bigoplus_{k\geq 1} \mathbf{Z}$ be a free $\mathbf{Z}$-module of countable rank. Then $M^\vee = \mathrm{Hom}_{\mathbf{Z}}(M, \mathbf{Z})$ is isomorphic to the direct product of countably many copies of $\mathbf{Z}$ by Theorem 3.3. We will show that $M^\vee$ is not a free $\mathbf{Z}$-module, so in contrast to Theorem 2.4 the dual of a free module of *infinite rank* may not be free.

Suppose, to the contrary, that a countable direct product $\prod_{k\geq 1} \mathbf{Z}$ is a free $\mathbf{Z}$-module. From the theory of modules over a PID, any submodule of a free module over a PID is free [1, pp. 650–651]. (The proof uses the well-ordering of a general set, which is logically equivalent to Zorn's lemma.) Consider the $\mathbf{Z}$-submodule $N \subset \prod_{k\geq 1} \mathbf{Z}$ consisting of integer sequences $(a_1, a_2, a_3, \dots)$ such that the highest power of 2 dividing a_k tends to ∞ as $k \to \infty$. For example, the sequences $a_k = 2^k$ and $a_k = k!$ are in N, as is any sequence where $a_k = 0$ for all large k. A sequence not in N is $a_k = k$. We are going to show N is not free. Therefore $\prod_{k\geq 1} \mathbf{Z}$ is not free either.

Let $e_i \in N$ be the vector with ith coordinate 1 and coordinate 0 elsewhere. These definitely are not a basis of N, since they aren't even a spanning set (example?). However, every element of $N/2N$ has only a finite number of nonzero coordinates, so the reductions $\overline{e}_i$ in $N/2N$ are a spanning set over $\mathbf{Z}/2\mathbf{Z}$. They are linearly independent over $\mathbf{Z}/2\mathbf{Z}$ too (check!), so the $\overline{e}_i$'s are a basis of $N/2N$. Thus $N/2N$ has countable dimension over $\mathbf{Z}/2\mathbf{Z}$.

If N were free, let $\{\alpha_i\}_{i\in I}$ be a $\mathbf{Z}$-basis of N. Then $N = \bigoplus_{i\in I} \mathbf{Z}\alpha_i$, so $N/2N = \bigoplus_{i\in I} (\mathbf{Z}/2\mathbf{Z})\overline{\alpha}_i$. Because we already checked $N/2N$ has countable dimension over $\mathbf{Z}/2\mathbf{Z}$, I must be countable. Then N has a countable basis and a countable scalar ring $\mathbf{Z}$, so N is countable.

On the other hand, the function $\prod_{k\geq 1} \mathbf{Z} \to N$ by $(a_1, a_2, a_3, \dots) \mapsto (2a_1, 4a_2, 8a_3, \dots)$ is injective, so N is uncountable because $\prod_{k\geq 1} \mathbf{Z}$ is uncountable. This is a contradiction, so N is not free, so $\prod_{k\geq 1} \mathbf{Z}$ is not free.

Dualizing turns direct sums into direct products. What can we say about the dual of a direct product? Is it isomorphic to the direct sum of the duals? In the next theorem we will write down an injective R-linear map in one direction, but there will not be an isomorphism in general when our direct products run over *infinite* index sets (which is the case when direct sums and direct products are different).

Theorem 3.6. *Let $M_i (i \in I)$ be R-modules. There is an injective R-linear map $\bigoplus_{i\in I} M_i^\vee \to (\prod_{i\in I} M_i)^\vee$ from the direct sum of the dual modules to the dual of the direct product module.*

Proof. Let $(\varphi_i)_{i\in I} \in \bigoplus_{i\in I} M_i^\vee$, so $\varphi_i \in M_i^\vee$ and all but finitely many φ_i's are zero maps. Then we can use these φ_i's to write down an R-linear map φ on the direct product:

$$\varphi((m_i)_{i\in I}) = \sum_{i\in I} \varphi_i(m_i).$$

Since all but finitely many φ_i's are zero maps, this sum is really only a finite sum (depending only on the m_i's for those i such that φ_i is not identically zero). The reader can check φ is R-linear and sending (φ_i) to φ is an R-linear map from $\bigoplus_{i\in I} M_i^\vee$ to $(\prod_{i\in I} M_i)^\vee$. We can recover the φ_i's from φ since $\varphi_i(m_i) = \varphi(m)$ where m has j-th coordinate 0 when $j \neq i$ and i-th coordinate m_i. Thus our map $\bigoplus_{i\in I} M_i^\vee \to (\prod_{i\in I} M_i)^\vee$ is injective. □

What is the image of the map we constructed in Theorem 3.6? For a $\varphi \in (\prod_{i\in I} M_i)^\vee$ which comes from a $(\varphi_i)_{i\in I} \in \bigoplus_{i\in I} M_i^\vee$, each value $\varphi((m_i)_{i\in I})$ depends only on the finitely many coordinates m_i at which φ_i is not identically 0. Let $\mathcal{F} = \{i : \varphi_i \neq 0\}$, a finite set, and take $N = \prod_{i\notin\mathcal{F}} M_i \times \prod_{i\in\mathcal{F}} \{0\}$, so $N \subset \ker\varphi$ and φ really lives on $(\prod_{i\in I} M_i)/N \cong \prod_{i\in\mathcal{F}} M_i$, a finite direct product. Conversely, any linear map $L\colon \prod_{i\in\mathcal{F}} M_i \to R$ can be lifted to an element of $(\prod_{i\in I} M_i)^\vee$ which kills N by projecting from the whole direct product onto the finite direct product $\prod_{i\in\mathcal{F}} M_i$ and then applying L. Thus, the image of the map in Theorem 3.6 is the elements of $(\prod_{i\in I} M_i)^\vee$ which depend on a finite set of coordinates. Reasonably, when I is infinite we can expect that most elements of $(\prod_{i\in I} M_i)^\vee$ do not depend on only finitely many coordinates, so our map in Theorem 3.6 should be very far from surjective. That is, it seems that the dual of a direct product is typically much larger than the direct sum of the dual modules, which we have naturally embedded as a submodule.

But perhaps we simply haven't been clever enough: might there be a different R-linear map $\bigoplus_{i\in I} M_i^\vee \to (\prod_{i\in I} M_i)^\vee$ which is a bijection? In general, no. We will see an example in Section 5 (Example 5.16).

4. Dual Bases and Double Duality

We return to the case of a finite free R-module M with rank $n > 0$. In Theorem 2.4 we saw that $M^\vee$ is finite free with rank n The proof came from stringing together the isomorphisms $M^\vee \cong (R^n)^\vee \cong R^n$. Let's look more closely at how to directly view $M^\vee$ as R^n. If we choose an R-basis $e_1, \ldots, e_n$ of M, then every $\varphi \in M^\vee$ is completely determined by its values on the e_i's, and sending φ to the n-tuple $(\varphi(e_1), \ldots, \varphi(e_n)) \in R^n$ is an R-linear injection from $M^\vee$ to R^n. It is surjective too, because each standard basis vector of R^n arises in this way (so the image contains the span of the standard basis, which is R^n). To see this, for $i = 1, 2, \ldots, n$ define $\varphi_i\colon M \to R$ by

$$\varphi_i(c_1e_1 + \cdots + c_ne_n) = c_i.$$

This is the ith coordinate function relative to our choice of basis of M. Under our map $M^\vee \to R^n$ this coordinate function goes to the ith standard basis vector of R^n. So we have

an isomorphism $M^\vee \to R^n$, and in this isomorphism the coordinate functions for our basis of M must be a basis of $M^\vee$ because the isomorphism identifies them with the standard basis of R^n.

So to each basis $e_1, \dots, e_n$ of a finite free R-module, the coordinate functions for this basis are a basis of the dual module. It is called the *dual basis* and is denoted $e_1^\vee, \dots, e_n^\vee$ (so φ_i above is $e_i^\vee$). They are the R-linear maps $M \to R$ determined by the conditions

$$e_i^\vee(e_j) = \begin{cases} 1, & \text{if } i = j, \\ 0, & \text{if } i \neq j. \end{cases} \tag{4.1}$$

Example 4.1. Let $R = \mathbf{R}$ and $M = \mathbf{C}$. What is the dual basis $\{f_1, f_2\}$ for $\mathbf{C}^\vee$ of the basis $\{1, i\}$ for $\mathbf{C}$? They are the coordinate functions for the basis $\{1, i\}$, so

$$f_1(a + bi) = a, \quad f_2(a + bi) = b$$

for real a and b. So $f_1 = \mathrm{Re}$ is the real part function and $f_2 = \mathrm{Im}$ is the imaginary part function. The basis of $\mathbf{C}^\vee$ which is dual to the basis $\{1, i\}$ of $\mathbf{C}$ is $\{\mathrm{Re}, \mathrm{Im}\}$.

While a finite free R-module M is isomorphic to its dual module $M^\vee$ by using bases, the isomorphism is not in any way canonical since a free module has no distinguished basis. However, if we work with the double-dual module $M^{\vee\vee} = (M^\vee)^\vee$ then there is a natural isomorphism with M in the finite free case.

Theorem 4.2. *When M is a finite-free R-module, there is a natural isomorphism $M \cong M^{\vee\vee}$.*

Proof. We will write down an R-linear map $M \to M^{\vee\vee}$ for *all* R-modules M, and then check it is an isomorphism when M is finite and free.

An element of $M^{\vee\vee}$ is a linear map $M^\vee \to R$. For each $m \in M$, evaluating elements of $M^\vee$ at m sends $M^\vee$ to R and is linear:

$$(\varphi + \psi)(m) = \varphi(m) + \psi(m), \quad (c\varphi)(m) = c(\varphi(m))$$

by the very definition of addition and scaling in $M^\vee$. Let $\mathrm{ev}_m \colon \varphi \mapsto \varphi(m)$ be this evaluation map, so $\mathrm{ev}_m \in M^{\vee\vee}$. Send M to $M^{\vee\vee}$ by $m \mapsto \mathrm{ev}_m$. This is additive since

$$\mathrm{ev}_{m+m'}(\varphi) = \varphi(m + m') = \varphi(m) + \varphi(m') = (\mathrm{ev}_m + \mathrm{ev}_{m'})(\varphi),$$

so $\mathrm{ev}_{m+m'} = \mathrm{ev}_m + \mathrm{ev}_{m'}$. (Notice that it was important that elements of $M^\vee$ are additive functions and not arbitrary functions from M to R.) Similarly, $\mathrm{ev}_{cm} = c\,\mathrm{ev}_m$ for $c \in R$. So sending m to ev_m is an R-linear map $M \to M^{\vee\vee}$ for all R-modules M.

We will now show that our map from M to $M^{\vee\vee}$ is an isomorphism when M is finite and free over R.

Let $e_1, \dots, e_n$ be an R-basis of M. Let $e_1^\vee, \dots, e_n^\vee$ be the dual basis of $M^\vee$. If $m \in M$ and $m \neq 0$, then m has a non-zero coordinate relative to our basis, so $e_i^\vee(m) \neq 0$ for some i. Therefore $\mathrm{ev}_m(e_i^\vee) \neq 0$, so ev_m is not the zero element of $M^{\vee\vee}$. This shows the only m for which ev_m is zero in $M^{\vee\vee}$ is 0, which means our map $M \to M^{\vee\vee}$ is injective.

It remains to show every element of $M^{\vee\vee}$ is some ev_m. We will use the dual basis for this. Pick an $f \in M^{\vee\vee}$. We want to find an $m \in M$ such that $f = \mathrm{ev}_m$, *i.e.*,

$$f(\varphi) = \varphi(m)$$

for all $\varphi \in M^\vee$. Since both sides of this equation are linear in φ, it suffices to find an m that makes this equation hold when φ runs through the dual basis $e_1^\vee, \dots, e_n^\vee$ since they

span $M^\vee$. Let $a_i = f(e_i^\vee) \in R$ and define $m = \sum_{i=1}^n a_i e_i \in M$. Then $f(e_i^\vee) = a_i = e_i^\vee(m)$, so $f = \mathrm{ev}_m$. □

The map $m \mapsto \mathrm{ev}_m$ we have constructed from M to $M^{\vee\vee}$ makes sense and is linear for all modules M, not just finite free modules, and will be called the *natural map* from a module to its double dual. For a finite free module M, this map is an isomorphism and will be called the *double duality isomorphism*. Under this isomorphism, the basis in $M^{\vee\vee}$ which is dual to the dual basis $e_1^\vee, \ldots, e_n^\vee$ in $M^\vee$ is the original basis $e_1, \ldots, e_n$. Indeed, as a parallel with (4.1) we have

$$\mathrm{ev}_{e_i}(e_j^\vee) = e_j^\vee(e_i) = \begin{cases} 1, & \text{if } i = j, \\ 0, & \text{if } i \neq j. \end{cases}.$$

Example 4.3. Here is an example of a non-free finitely generated module for which the natural map to the double dual is *not* an isomorphism. Take $R = \mathbf{Z}[X]$ and $M = (2, X)$. From Example 2.9, $M^\vee \cong R$. Therefore $M^{\vee\vee} \cong R^\vee \cong R$, so M is not isomorphic to its (dual or) double dual since M is a non-principal ideal in R.

Example 4.4. So you don't get the idea that the double duality isomorphism $M \cong M^{\vee\vee}$ works *only* for finite-free modules, let's discuss a case where this isomorphism occurs and M is finitely generated but not free. Consider the ideals $I = (3, 1 + \sqrt{-14})$ and $J = (3, 1 - \sqrt{-14})$ in the ring $R = \mathbf{Z}[\sqrt{-14}]$. From Example 2.8 each of these ideals can be viewed as the dual R-module of the other ideal by making $x \in I$ and $y \in J$ act on each other as $\frac{xy}{3} \in R$. From the way I and J act as each other's dual module, the natural maps $I \to I^{\vee\vee}$ and $J \to J^{\vee\vee}$ are both isomorphisms (check!).

The R-modules I and J are not free. We give the argument only for I; a similar argument will work for J. (What we are about to do is tedious algebra. There are more efficient approaches using algebraic number theory.) For any x and x' in I we have the R-linear relation $cx + c'x' = 0$ with $c = x'$ and $c' = -x$, so a linearly independent subset of I has size at most 1: if I is free then I has a one-element basis $\{\alpha\}$, meaning $I = R\alpha$ is a principal ideal. We will show I is not principal, by contradiction. Suppose $I = (\alpha)$ and write $\alpha = a + b\sqrt{-14}$ with integers a and b. Since $3 \in I$, $3 = (a + b\sqrt{-14})(a' + b'\sqrt{-14})$ for some $a' + b'\sqrt{-14}$ in $\mathbf{Z}[\sqrt{-14}]$. Taking complex absolute values of both sides and squaring, $9 = (a^2 + 14b^2)(a'^2 + 14b'^2)$. Since $a^2 + 14b^2$ is a non-negative integer dividing 9, it can only be 1, 3, or 9. The first choice implies $a + b\sqrt{-14} = \pm 1$, so $I = (\pm 1) = (1)$ is the unit ideal. The second choice is impossible ($a^2 + 14b^2 = 3$ has no integral solution). The third choice implies $a + b\sqrt{-14} = \pm 3$ so $I = (3)$. Since $J = (3, 1 - \sqrt{-14})$ is the complex-conjugate ideal to I (that is, taking complex conjugates of the elements in an ideal turns I into J, and *vice versa*), if $I = (1)$ then $J = (1)$ while if $I = (3)$ then $J = (3)$. Neither of these is compatible with $IJ = (3)$, so I is not a principal ideal and thus is not a free R-module.

This is not only an example where a non-free module (either I or J) is isomorphic to its double dual by the natural map, but also the module is *not* isomorphic to its dual. While *finite free* modules are non-canonically isomorphic to their duals, we are saying $I \ncong J$ as R-modules. (Notice $I \cong J$ as $\mathbf{Z}$-modules, since complex conjugation is an additive isomorphism, but it is not R-linear!) To show $I \ncong J$ as R-modules, suppose there is an R-module isomorphism $\varphi\colon I \to J$. For any $x \in I$, $\varphi(3x) = 3\varphi(x) = x\varphi(3)$, so $\varphi(x) = (\varphi(3)/3)x$. This tells us φ is a scaling map from I to J by the (mysterious) factor $\varphi(3)/3$. Therefore $J = \varphi(I) = \frac{\varphi(3)}{3} I$, so $3J = \varphi(3)I$. Multiplying both sides by the ideal I, $3IJ = \varphi(3)I^2$, so $(9) = \varphi(3)I^2$ (because $IJ = (3)$). Therefore I^2 is a principal ideal.

However, $I^2 = (9, 2 - \sqrt{-14})$ and $I^4 = (5 + 2\sqrt{-14})$ (check!), so if I^2 were principal, say $I^2 = (\alpha)$, then $\alpha^2 = \pm(5 + 2\sqrt{-14})$. But $\pm(5 + 2\sqrt{-14})$ is not a square in R, so we have a contradiction.

Remark 4.5. Example 4.4 is not an isolated curiosity. There are many finitely generated modules which are isomorphic to their double duals by the natural map and are not free modules. For instance, if R is the ring of algebraic integers of a number field then any ideal $I \subset R$ is isomorphic to its double dual by the natural map, and I is not free when I is a non-principal ideal. Moreover, I is not isomorphic to its dual module if the order of I in the ideal class group of R is greater than 2. Example 4.4 is a special case of this: $\mathbf{Z}[\sqrt{-14}]$ is the ring of algebraic integers of $\mathbf{Q}(\sqrt{-14})$ and the ideal $(3, 1 + \sqrt{-14})$ has order 4 in the ideal class group.

Example 4.6. In Example 3.5 we saw that if M is a countable direct sum of copies of $\mathbf{Z}$ then its $\mathbf{Z}$-dual is not free. But the natural map $M \to M^{\vee\vee}$ is nevertheless an isomorphism. We omit the proof.

Modules for which the natural map $M \to M^{\vee\vee}$ is an isomorphism are called *reflexive*. Thus finite free modules are reflexive, as are the modules in Examples 4.4 and 4.6.

Example 4.7. Let V be any infinite-dimensional vector space over a field K. Using Zorn's lemma one can show $\dim_K(V) < \dim_K(V^\vee) < \dim_K(V^{\vee\vee})$, so V is not isomorphic to its double dual. A special case of this will be discussed in Example 5.16.

Remark 4.8. In analysis, one often deals with infinite-dimensional vector spaces V over $\mathbf{R}$ or $\mathbf{C}$. Usually there is a topology on V lurking around and one *redefines* $V^\vee$ to denote the continuous linear maps from V to $\mathbf{R}$ (or $\mathbf{C}$), rather than all linear maps to the scalars. This "continuous" dual space is much smaller than the "algebraic" dual space, and can be equipped with its own topology. In many cases with this new definition of the dual space, the natural map $V \to V^{\vee\vee}$ is an isomorphism just like in the finite-dimensional setting. Here we are doing everything algebraically; there are no topologies on our modules.

5. Dual Maps

Let M and N be R-modules and $L\colon M \to N$ be R-linear. We can use L to turn functionals on N into functionals on M: if $\varphi \in N^\vee$ then $\varphi \circ L \in M^\vee$. So there is a map $L^\vee\colon N^\vee \to M^\vee$ given by the rule $L^\vee(\varphi) = \varphi \circ L$. Notice $L\colon M \to N$ and $L^\vee\colon N^\vee \to M^\vee$. Here is a picture:

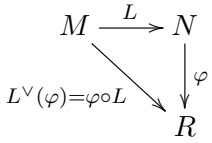

We call $L^\vee$ the *dual map* to L.

Example 5.1. Let $R = \mathbf{R}$, $M = \mathbf{R}[X]$, $N = \mathbf{R}^2$ and $L\colon M \to N$ by $L(f(X)) = (f(0), f(1))$. The map $\varphi\colon \mathbf{R}^2 \to \mathbf{R}$ given by $\varphi(x, y) = 2x+3y$ is in the dual space of $\mathbf{R}^2$ and the composite $\varphi \circ L$ which sends $f(X)$ to $2f(0) + 3f(1)$ is in the dual space of $\mathbf{R}[X]$. This composite is $L^\vee(\varphi)$.

Example 5.2. Let $L\colon M \to N$ be the zero map: $L(m) = 0$ for all $m \in M$. Then $L^\vee\colon N^\vee \to M^\vee$ satisfies $(L^\vee(\varphi))(m) = \varphi(L(m)) = \varphi(0) = 0$ for all $m \in M$ and $\varphi \in N^\vee$, so $L^\vee$ is the zero map on the dual modules.

Theorem 5.3. *The map $L^\vee\colon N^\vee \to M^\vee$ is R-linear.*

Proof. For φ and ψ in $N^\vee$ and m in M,

$$\begin{aligned}(L^\vee(\varphi+\psi))(m) &= (\varphi+\psi)(L(m))\\ &= \varphi(L(m))+\psi(L(m))\\ &= (L^\vee(\varphi))(m)+(L^\vee(\psi))(m)\\ &= (L^\vee(\varphi)+L^\vee(\psi))(m).\end{aligned}$$

Since this holds for all $m \in M$, $L^\vee(\varphi+\psi) = L^\vee(\varphi)+L^\vee(\psi)$ in $M^\vee$.

For φ in $N^\vee$, $c \in R$, and $m \in M$,

$$\begin{aligned}(L^\vee(c\varphi))(m) &= (c\varphi)(L(m))\\ &= c(\varphi(L(m)))\\ &= c((L^\vee(\varphi))(m))\\ &= (c(L^\vee(\varphi)))(m).\end{aligned}$$

Since this holds for all $m \in M$,

$$L^\vee(c\varphi) = c(L^\vee(\varphi)).$$

□

Theorem 5.4. *The function $\mathrm{Hom}_R(M,N) \to \mathrm{Hom}_R(N^\vee, M^\vee)$ given by $L \mapsto L^\vee$ is R-linear.*

Proof. For L_1 and L_2 in $\mathrm{Hom}_R(M,N)$ we want $(L_1+L_2)^\vee = L_1^\vee + L_2^\vee$, which means $(L_1+L_2)^\vee(\varphi) = (L_1^\vee + L_2^\vee)(\varphi)$ for all $\varphi \in N^\vee$. Both sides are in $M^\vee$, so we check equality by evaluating both sides separately at any $m \in M$:

$$\begin{aligned}((L_1+L_2)^\vee(\varphi))(m) &= (\varphi\circ(L_1+L_2))(m)\\ &= \varphi((L_1+L_2)(m))\\ &= \varphi(L_1(m)+L_2(m))\\ &= \varphi(L_1(m))+\varphi(L_2(m))\end{aligned}$$

since φ is additive. Also

$$\begin{aligned}((L_1^\vee+L_2^\vee)(\varphi))(m) &= (L_1^\vee(\varphi)+L_2^\vee(\varphi))(m)\\ &= (\varphi\circ L_1+\varphi\circ L_2)(m)\\ &= \varphi(L_1(m))+\varphi(L_2(m)).\end{aligned}$$

These agree for all m, so $(L_1+L_2)^\vee(\varphi) = (L_1^\vee+L_2^\vee)(\varphi)$ in $M^\vee$. Therefore $(L_1+L_2)^\vee$ and $L_1^\vee+L_2^\vee$ have the same value at every $\varphi \in N^\vee$, so $(L_1+L_2)^\vee = L_1^\vee+L_2^\vee$ in $\mathrm{Hom}_R(N^\vee, M^\vee)$.

The proof that $(cL)^\vee = cL^\vee$ for all $c \in R$ and $L \in \mathrm{Hom}_R(M,N)$ is left to the reader. □

Example 5.5. The dual of the identity map on M is the identity map on $M^\vee$. Indeed, for $\varphi \in M^\vee$,

$$\mathrm{id}_M^\vee(\varphi) = \varphi\circ\mathrm{id}_M = \varphi,$$

so $\mathrm{id}_M^\vee = \mathrm{id}_{M^\vee}$. Similarly, for any $c \in R$, $(c\,\mathrm{id}_M)^\vee = c\,\mathrm{id}_{M^\vee}$.

When M and N are finite free non-zero R-modules, picking bases lets us turn the module $\mathrm{Hom}_R(M,N)$ of linear maps between M and N into matrices. The dual modules $M^\vee$ and $N^\vee$ are finite free so the module $\mathrm{Hom}_R(N^\vee, M^\vee)$ can also be turned into matrices once

bases of the dual modules are chosen. If the bases we use for $M^\vee$ and $N^\vee$ are the dual bases to the bases chosen for M and N, then there is a close relation between the matrix representations of any L in $\mathrm{Hom}_R(M,N)$ and its dual map $L^\vee$. We will compute an example before treating the general situation.

Example 5.6. Let $R = \mathbf{R}$ and $M = N = \mathbf{C}$. Set $L\colon \mathbf{C} \to \mathbf{C}$ by $L(z) = (2+i)z + \overline{z}$. This is $\mathbf{R}$-linear, with

$$L(1) = 3 + i, \quad L(i) = -1 + i.$$

Therefore, relative to the basis $\{1, i\}$ of $\mathbf{C}$ (for inputs and outputs), L is represented by the matrix $\left(\begin{smallmatrix} 3 & -1 \\ 1 & 1 \end{smallmatrix}\right)$. What is the matrix for $L^\vee$ relative to the dual basis of $\mathbf{C}^\vee$?

We saw in Example 4.1 that the dual basis to $\{1, i\}$ is $\{\mathrm{Re}, \mathrm{Im}\}$. We want to express $L^\vee(\mathrm{Re})$ and $L^\vee(\mathrm{Im})$ in terms of the dual basis. Note $L^\vee(\mathrm{Re}) = \mathrm{Re} \circ L$ and $L^\vee(\mathrm{Im}) = \mathrm{Im} \circ L$. That is, we want to compute the real and imaginary parts of the values of L. For every $z = a + bi$ in $\mathbf{C}$,

$$L(z) = (2+i)z + \overline{z} = (2+i)(a+bi) + a - bi = (3a - b) + (a+b)i. \tag{5.1}$$

Therefore $L(z)$ has real part $3a - b = 3\,\mathrm{Re}(z) - \mathrm{Im}(z)$ and $L(z)$ has imaginary part $a + b = \mathrm{Re}(z) + \mathrm{Im}(z)$, which means

$$L^\vee(\mathrm{Re}) = 3\,\mathrm{Re} - \mathrm{Im}, \quad L^\vee(\mathrm{Im}) = \mathrm{Re} + \mathrm{Im}\,.$$

The matrix for $L^\vee$ in the basis $\{\mathrm{Re}, \mathrm{Im}\}$ is $\left(\begin{smallmatrix} 3 & 1 \\ -1 & 1 \end{smallmatrix}\right)$. This matrix is the transpose of the matrix found for L relative to the basis $\{1, i\}$.

What we found in this example is true in general: the matrix representations of L and $L^\vee$ are transposes of each other when using any choice of bases for M and N and the *dual* bases to those choices in $M^\vee$ and $N^\vee$. The proof of this result, which follows below, may seem rather abstract, but it's really just a matter of unwinding lots of definitions. Before we begin, at least notice that when M has rank m and N has rank n, $\mathrm{Hom}_R(M,N)$ becomes $n \times m$ matrices and $\mathrm{Hom}_R(N^\vee, M^\vee)$ becomes $m \times n$ matrices when using bases, so the row and column sizes get flipped. Thus it is no surprise that matrix transposes can get involved.

Theorem 5.7. *Let M and N be non-zero finite free R-modules. Pick bases $\mathcal{B} = \{e_1, \ldots, e_m\}$ of M and $\mathcal{C} = \{f_1, \ldots, f_n\}$ of N. Let their dual bases of $M^\vee$ and $N^\vee$ be denoted $\mathcal{B}^\vee = \{e_1^\vee, \ldots, e_m^\vee\}$ and $\mathcal{C}^\vee = \{f_1^\vee, \ldots, f_n^\vee\}$. For a linear map $L\colon M \to N$, the matrices ${}_{\mathcal{C}}[L]_{\mathcal{B}}$ and ${}_{\mathcal{B}^\vee}[L^\vee]_{\mathcal{C}^\vee}$ are transposes.*

Proof. Set $A = {}_{\mathcal{C}}[L]_{\mathcal{B}}$ and $A' = {}_{\mathcal{B}^\vee}[L^\vee]_{\mathcal{C}^\vee}$. These are matrices: A is $n \times m$ and A' is $m \times n$. We want to show $A' = A^\top$.

Let $[\cdot]_{\mathcal{B}}\colon M \to R^m$ and $[\cdot]_{\mathcal{C}}\colon N \to R^n$ be the R-module isomorphisms coming from coordinates using bases $\mathcal{B}$ and $\mathcal{C}$. Similarly define $[\cdot]_{\mathcal{B}^\vee}\colon M^\vee \to R^m$ and $[\cdot]_{\mathcal{C}^\vee}\colon N^\vee \to R$.

The matrices A and A' realize the transformations L and $L^\vee$ in the chosen bases. That is,

$$[L(m)]_{\mathcal{C}} = A[m]_{\mathcal{B}}, \quad [L^\vee(\varphi)]_{\mathcal{B}^\vee} = A'[\varphi]_{\mathcal{C}^\vee} \tag{5.2}$$

for $m \in M$ and $\varphi \in N^\vee$.

What is the (i,j) entry of A? Since $[e_j]_{\mathcal{B}}$ is the j-th standard basis vector of R^m, the j-th column of A is $A[e_j]_{\mathcal{B}} = [L(e_j)]_{\mathcal{C}}$, which is the coordinate vector of $L(e_j)$ in the chosen basis $\mathcal{C}$ for N. The i-th coordinate of the vector $[L(e_j)]_{\mathcal{C}} \in R^n$ is $f_i^\vee(L(e_j))$ since f_i is the i-th basis vector in $\mathcal{C}$.

Now we compute the (j,i) entry of A'. The i-th column of A' is $A'[f_i^\vee]_{\mathcal{C}^\vee}$ since $[f_i^\vee]_{\mathcal{C}^\vee}$ is the i-th standard basis vector of R^n. By (5.2), $A'[f_i^\vee]_{\mathcal{C}^\vee} = [L^\vee(f_i^\vee)]_{\mathcal{B}^\vee}$. The coordinate functions on $M^\vee$ relative to the basis $\mathcal{B}^\vee$ are the dual basis to this dual basis, which means they are the original basis $\mathcal{B}$ when we identify M with $M^{\vee\vee}$ by the natural map. So the j-th coordinate of the vector $[L^\vee(f_i^\vee)]_{\mathcal{B}^\vee}$ is $\mathrm{ev}_{e_j}(L^\vee(f_i^\vee)) = (L^\vee(f_i^\vee))(e_j)$. By the definition of $L^\vee$, $L^\vee(f_i^\vee) = f_i^\vee \circ L$, so

$$(L^\vee(f_i^\vee))(e_j) = f_i^\vee(L(e_j)).$$

This is what we computed in the previous paragraph as the (i,j) entry of A. Thus A and A' are transposed matrices. □

Up to now we have dealt with the dual of an individual linear map. We can also regard dualizing linear maps as a construction on all of $\mathrm{Hom}_R(M,N)$ at once: it is a function $\mathrm{Hom}_R(M,N) \xrightarrow{\vee} \mathrm{Hom}_R(N^\vee, M^\vee)$ between spaces of linear maps. From this point of view, we can express Theorem 5.7 as a commutative diagram

$$\begin{array}{ccc} \mathrm{Hom}_R(M,N) & \xrightarrow{(\cdot)^\vee} & \mathrm{Hom}_R(N^\vee, M^\vee) \\ {\scriptstyle \mathcal{C}[\cdot]_{\mathcal{B}}}\downarrow & & \downarrow{\scriptstyle \mathcal{B}^\vee[\cdot]_{\mathcal{C}^\vee}} \\ \mathrm{M}_{n\times m}(R) & \xrightarrow[(\cdot)^\top]{} & \mathrm{M}_{m\times n}(R), \end{array}$$

where the side maps are R-module isomorphisms which come from choosing coordinates relative to a choice of bases and dual bases, the top map is the dual construction for linear maps, and the bottom map is the matrix transpose. Since the side maps are isomorphisms, all properties of the matrix transpose can be transferred to the dual map for finite free R-modules. But that is the wrong way to think. We should instead reprove all familiar properties of the matrix transpose directly in terms of the dual map construction when possible. This is important because the dual map makes sense even when modules are not finite free. The matrix transpose is merely the *special case* of the dual map construction for finite free R-modules.

For instance, Theorem 5.4 generalizes the linearity of the matrix transpose ($(A+B)^\top = A^\top + B^\top, (cA)^\top = cA^\top$). Example 5.5 generalizes the fact that a scalar diagonal matrix is its own transpose. The following theorem generalizes the reverse multiplicativity $(AB)^\top = B^\top A^\top$ for the tranpose of matrix products. It *conceptually* explains why the matrix transpose reverses the order of multiplication, in much the same way that interpreting matrix multiplication as composition of transformations explains associativity and non-commutativity of matrix products from associativity and non-commutativity of function composition.

Theorem 5.8. *Let M, N, and P be any R-modules. If $L_1 \colon M \to N$ and $L_2 \colon N \to P$ are linear, then the composite $L_2 \circ L_1 \colon M \to P$ has dual*

$$(L_2 \circ L_1)^\vee = L_1^\vee \circ L_2^\vee.$$

Proof. Both maps go from $P^\vee$ to $M^\vee$. For $\varphi \in P^\vee$ and $m \in M$,

$$\begin{aligned}((L_2 \circ L_1)^\vee(\varphi))(m) &= (\varphi(L_2 \circ L_1))(m)\\ &= \varphi(L_2(L_1(m)))\\ &= (L_2^\vee(\varphi))(L_1(m))\\ &= (L_1^\vee(L_2^\vee(\varphi)))(m)\\ &= ((L_1^\vee \circ L_2^\vee)(\varphi))(m).\end{aligned}$$

Since we have equality at all m,

$$(L_2 \circ L_1)^\vee(\varphi) = (L_1^\vee \circ L_2^\vee)(\varphi)$$

for all $\varphi \in P^\vee$, so $(L_2 \circ L_1)^\vee = L_1^\vee \circ L_2^\vee$ in $\mathrm{Hom}_R(P^\vee, M^\vee)$. □

Example 5.5 and Theorem 5.8 are called the *functoriality* of the dual map construction. It means the dual of an identity map is an identity map and dualizing interacts with composition in a definite manner (reversing the order of composition, actually).

Corollary 5.9. *When $L\colon M \to N$ is an isomorphism of R-modules, $L^\vee$ is an isomorphism of the dual modules and $(L^\vee)^{-1} = (L^{-1})^\vee$.*

Proof. We have $L^{-1} \circ L = \mathrm{id}_M$ and $L \circ L^{-1} = \mathrm{id}_N$. Passing to the dual maps on both sides and using Example 5.5 and Theorem 5.8, the composites of $L^\vee$ and $(L^{-1})^\vee$ in both directions are the identity maps on $M^\vee$ and $N^\vee$. □

This generalizes the identity $(A^{-1})^\top = (A^\top)^{-1}$ when A is an invertible square matrix.

Corollary 5.10. *For finite free R-modules M and N, $L \mapsto L^\vee$ is an R-module isomorphism from $\mathrm{Hom}_R(M,N)$ to $\mathrm{Hom}_R(N^\vee, M^\vee)$.*

Proof. Since M and N are finite free, so are their dual modules, and thus $\mathrm{Hom}_R(M,N)$ and $\mathrm{Hom}_R(N^\vee, M^\vee)$ are finite free. We will show dualizing take a basis of $\mathrm{Hom}_R(M,N)$ to a basis of $\mathrm{Hom}_R(N^\vee, M^\vee)$ and therefore is an isomorphism.

Pick bases $\{e_1, \ldots, e_m\}$ of M and $\{f_1, \ldots, f_n\}$ of N. A basis of $\mathrm{Hom}_R(M,N)$ is the functions $L_{ij}\colon M \to N$ where $L_{ij}(e_i) = f_j$ and $L_{ij}(e_k) = 0$ for $k \neq i$:

$$L_{ij}(e_k) = \begin{cases} f_j, & \text{if } k = i,\\ 0, & \text{if } k \neq i.\end{cases}$$

That is, $L_{ij}(a_1e_1 + \cdots + a_re_r) = a_if_j$. We will show the dual maps $L_{ij}^\vee\colon N^\vee \to M^\vee$ form a similar type of basis for $\mathrm{Hom}_R(N^\vee, M^\vee)$.

What is the effect of $L_{ij}^\vee$ on the dual basis $f_1^\vee, \ldots, f_s^\vee$? For any basis vector e_k,

$$\begin{aligned}(L_{ij}^\vee(f_\ell^\vee))(e_k) &= (f_\ell^\vee \circ L_j^\vee)(e_k)\\ &= f_\ell^\vee(L_{ij}(e_k))\\ &= \begin{cases} f_\ell^\vee(f_j), & \text{if } k = i,\\ 0, & \text{if } k \neq i,\end{cases}\\ &= \begin{cases} 1, & \text{if } k = i, \ell = j,\\ 0, & \text{otherwise,}\end{cases}\end{aligned}$$

so

$$L_{ij}^\vee(f_\ell^\vee) = \begin{cases} e_i^\vee, & \text{if } \ell = j, \\ 0, & \text{if } \ell \neq j. \end{cases}$$

Thus the functions $\{L_{ij}^\vee\}$ are a basis of $\mathrm{Hom}_R(N^\vee, M^\vee)$. □

If $M = N$ then $\mathrm{Hom}_R(M,N) = \mathrm{Hom}_R(M,M)$ and $\mathrm{Hom}_R(N^\vee, M^\vee) = \mathrm{Hom}_R(M^\vee, M^\vee)$, are both rings, and passing from one ring to the other using the dual map construction is a ring anti-homomorphism (that is, it's a homomorphism except the order of multiplication is reversed), which is a ring anti-isomorphism when M is finite free. It generalizes the matrix transpose on $\mathrm{M}_n(R)$.

With dual maps we can generalize the identity $A^{\top\top} = A$ for matrices A, as follows.

Theorem 5.11. *Let M and N be finite free R-modules. For any linear map $L\colon M \to N$, the double dual map $L^{\vee\vee}\colon N^{\vee\vee} \to M^{\vee\vee}$ is identified with L when we use the double duality isomorphism from Theorem 4.2. That is, the diagram*

$$\begin{array}{ccc} M & \xrightarrow{\quad L \quad} & N \\ \downarrow & & \downarrow \\ M^{\vee\vee} & \xrightarrow{\quad L^{\vee\vee} \quad} & N^{\vee\vee} \end{array}$$

commutes, where the vertical maps are the double duality isomorphisms.

Proof. Pick $m \in M$. Taking m down the left side turns it into ev_m in $M^{\vee\vee}$, and applying $L^{\vee\vee}$ to that gives us

$$L^{\vee\vee}(\mathrm{ev}_m) = (L^\vee)^\vee(\mathrm{ev}_m) = \mathrm{ev}_m \circ L^\vee$$

in $N^{\vee\vee}$. On the other hand, if we take m across the top and then down the right side we get $\mathrm{ev}_{L(m)}$. Are $\mathrm{ev}_m \circ L^\vee$ and $\mathrm{ev}_{L(m)}$ the same element of $N^{\vee\vee}$? Well, for any $\varphi \in N^\vee$ we have

$$(\mathrm{ev}_m \circ L^\vee)(\varphi) = \mathrm{ev}_m(L^\vee(\varphi)) = \mathrm{ev}_m(\varphi \circ L) = (\varphi \circ L)(m) = \varphi(L(m)),$$

which is the meaning of $\mathrm{ev}_{L(m)}(\varphi)$. So we get equality for all m in M, which means the diagram commutes. □

We used finite freeness of M and N in the proof of Theorem 5.11 to know the natural maps $M \to M^{\vee\vee}$ and $N \to N^{\vee\vee}$ are isomorphisms. So Theorem 5.11 is true for any R-modules where the natural map to the double dual is an isomorphism. Examples 4.4 and 4.6 give such examples that are not finite free.

What does duality do to injectivity and surjectivity of a linear map?

Theorem 5.12. *Let $L\colon M \to N$ be a linear map of R-modules. If L is onto then $L^\vee$ is one-to-one.*

Proof. Suppose $\varphi \in N^\vee$ and $L^\vee(\varphi) = 0$. We want to show $\varphi = 0$. Well, by the definition of the dual map $\varphi \circ L = 0$ as a function from M to R, so $\varphi(L(m)) = 0$ for all $m \in M$. Since L is onto, $\{L(m) : m \in M\} = N$. Thus $\varphi = 0$ as a function on N. □

So duality converts surjectivity into injectivity. Does it convert injectivity into surjectivity? Well, if $L\colon M \to N$ is one-to-one then we can use L to view M as a submodule of N: $M \cong L(M) \subset N$. For $\varphi \in N^\vee$, the map $L^\vee(\varphi) = \varphi \circ L$ is, from this point of view, simply the restriction of φ to the submodule $L(M) \subset N$. To say $L^\vee$ is onto means every linear map $\psi\colon M \to R$ has the form $\varphi \circ L$, which is saying (since $M \cong L(M)$ using L) that every

linear map $L(M) \to R$ can be extended to a linear map $N \to R$. So N has the property that all elements of the dual of the submodule $L(M) \subset N$ extend to elements of the dual of N. This property is not generally true! Here's an example.

Example 5.13. Let $R = \mathbf{Z}$, $M = 2\mathbf{Z}$, and $N = \mathbf{Z}$. Let $L\colon M \to N$ be the natural inclusion. We will show $L^\vee\colon N^\vee \to M^\vee$ is not onto. Set $\varphi\colon M \to R$ by $\varphi(2a) = a$, so $\varphi \in M^\vee$. Saying φ is in the image of $L^\vee$ means φ has an extension (using L) to $N^\vee$, say Φ, so $\Phi(2) = \varphi(2) = 1$ and $\Phi(2) = 2\Phi(1)$. Thus $2\Phi(1) = 1$, but this has no solution for $\Phi(1)$ in $R = \mathbf{Z}$. So Φ does not exist.

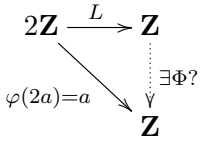

In fact, the image of $L^\vee$ is the elements of $M^\vee$ whose image in R is $2\mathbf{Z}$. A similar example can be made using $k\mathbf{Z}$ for M with any $k > 1$.

Despite this example, there is an important case when duality does turn injective linear maps into surjective linear maps: when R is a field. Let's prove this.

Theorem 5.14. *Let F be a field and M and N be F-vector spaces with $L\colon M \to N$ a linear map. If L is one-to-one then $L^\vee$ is onto.*

Proof. This is obvious for $M = 0$ or $N = 0$, so take M and N nonzero. Since we are working over a field, the subspace $L(M)$ of N has a direct sum complement: by choosing a basis of $L(M)$ and extending it to a basis of N we can write

$$N = L(M) \oplus P \tag{5.3}$$

for a subspace P. (Every F-vector space has a basis and a basis for a subspace can be enlarged to a basis of the whole space. Apply this result by starting with a basis for $L(M)$, extend it to a basis of N, and take for P the span of the additional part of the basis for N not coming from the basis for $L(M)$.) Every linear map $L(M) \to F$ can be extended to a linear map $N \to F$ by projecting from N to $L(M)$ by the direct sum decomposition (5.3) and then applying the chosen linear map on $L(M)$. Now the argument preceding Example 5.13 goes through to show $L^\vee$ is onto.

Here is another way of ending this argument. Let $\pi\colon N \to M$ be the projection from N to $L(M)$ followed by undoing L (which is possible since L is one-to-one from M onto $L(M)$). That is, $\pi(L(m) + p) = m$. This is linear and $\pi \circ L = \mathrm{id}_M$ (check!). Now dualize to get $L^\vee \circ \pi^\vee = \mathrm{id}_M^\vee = \mathrm{id}_{M^\vee}$. This implies $L^\vee$ is onto, since for any $\varphi \in M^\vee$ we have

$$L^\vee(\pi^\vee(\varphi)) = \varphi.$$

□

Theorem 5.14 applies not only to finite-dimensional vector spaces, but to all vector spaces: any nonzero vector space has a basis and a basis of a subspace always extends to a basis of the whole space. In the infinite-dimensional case that needs Zorn's Lemma, so it is definitely not constructive.

Corollary 5.15. *Let M and N be R-modules and $L\colon M \to N$ be a linear map. If L is one-to-one and $L(M)$ is a direct summand of N, then $L^\vee$ is onto.*

Proof. The hypothesis is that $N = L(M) \oplus P$ for some submodule P. This hypothesis matches (5.3), which was derived in the vector space case using bases. In the module case, by making this property a hypothesis we can make the proof of Theorem 5.14 go through for modules. □

Notice $2\mathbf{Z}$ is not a direct summand of $\mathbf{Z}$, thus "explaining" Example 5.13.

To see that Theorem 5.14 is useful in the infinite-dimensional setting, we now use it to show that the dual module of a direct product can be much larger than the direct sum of the dual modules (compare with Example 4.6).

Example 5.16. Let F be a finite field (such as $\mathbf{Z}/p\mathbf{Z}$) and set $V = \bigoplus_{n\geq 1} F$ to be a direct sum of countably many copies of F. This is countable and by Theorem 3.3, $V^\vee \cong \prod_{n\geq 1}(F^\vee)$, which is uncountable. By Theorem 5.14, the dual map to the inclusion $\bigoplus_{n\geq 1} F \hookrightarrow \prod_{n\geq 1} F$ is a surjective map on the dual spaces in the other direction:

$$\left(\prod_{n\geq 1} F\right)^\vee \twoheadrightarrow \left(\bigoplus_{n\geq 1} F\right)^\vee = V^\vee.$$

Therefore $(\prod_{n\geq 1} F)^\vee$ admits a surjection to an uncountable set, so $(\prod_{n\geq 1} F)^\vee$ is uncountable. Since $F^\vee \cong F$ as F-vector spaces, $V^{\vee\vee} \cong (\prod_{n\geq 1} F^\vee)^\vee \cong (\prod_{n\geq 1} F)^\vee$, so $V^{\vee\vee}$ is uncountable. This implies that the natural map $V \to V^{\vee\vee}$ (and in fact any linear map $V \to V^{\vee\vee}$) is not surjective.

We know (Example 5.13) that the dual of an injective linear map need not be a surjective linear map in general. This is the only way in which dualizing does not behave well on exact sequences of linear maps. Here is the basic result:

Theorem 5.17. *Dualizing is right exact: if* $M \xrightarrow{f} N \xrightarrow{g} P \xrightarrow{h} 0$ *is exact then the dual sequence* $0 \xrightarrow{h^\vee} P^\vee \xrightarrow{g^\vee} N^\vee \xrightarrow{f^\vee} M^\vee$ *is exact.*

Proof. Since $g \circ f$ and $h \circ g$ are both the map 0, dualizing implies $f^\vee \circ g^\vee$ and $g^\vee \circ h^\vee$ equal 0 (Example 5.2). So the image of $h^\vee$ is in the kernel of $g^\vee$ and the image of $g^\vee$ is in the kernel of $f^\vee$. It remains to show the kernel of $g^\vee$ is in the image of $h^\vee$ and the kernel of $f^\vee$ is in the image of $g^\vee$.

Since g is onto, by Theorem 5.12 $g^\vee$ is one-to-one, so the kernel of $g^\vee$ is 0 in $P^\vee$, which is also the image of $h^\vee$.

Next, suppose $\varphi \in N^\vee$ is in the kernel of $f^\vee$, so $f^\vee(\varphi) = \varphi \circ f$ is the zero map.

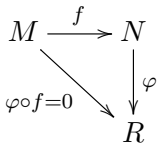

We want to show $\varphi = g^\vee(\psi) = \psi \circ g$ for some $\psi \in P^\vee$.

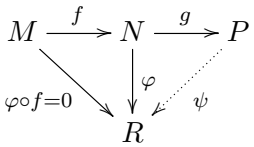

To construct such a ψ, we work backwards across the top of the diagram. For any $x \in P$, we can write $x = g(n)$ for some $n \in N$ since g is onto. If also $x = g(n')$ then $n - n' \in \ker g$, which is the image of f. Since $\varphi \circ f = 0$, φ is 0 on the image of f, so $\varphi(n - n') = 0$, so $\varphi(n) = \varphi(n')$. Thus defining $\psi\colon P \to R$ by $\psi(x) = \varphi(n)$ where $x = g(n)$ is well-defined. As an exercise, check ψ is linear. By construction, $\psi(g(n)) = \varphi(n)$, so $\psi \circ g = \varphi$. □

Example 5.18. A basic result in matrix algebra says the span of the columns of a matrix over a field has the same dimension as the span of the rows of the matrix: row rank equals column rank. We will use Theorem 5.17 to give a generalization of this theorem to *any* linear map between *any* two modules (even non-free modules) over *any* commutative ring.

Consider an R-linear map $L\colon M \to N$ between two R-modules M and N. It may not be surjective, but the reduction map $\pi_L\colon N \to N/L(M)$ is surjective. So we have the exact sequence

$$M \xrightarrow{L} N \xrightarrow{\pi_L} N/L(M) \longrightarrow 0.$$

Dualizing,

$$0 \longrightarrow (N/L(M))^\vee \xrightarrow{\pi_L^\vee} N^\vee \xrightarrow{L^\vee} M^\vee$$

is exact, so

$$\operatorname{im}(\pi_L^\vee) = \ker(L^\vee). \tag{5.4}$$

Let's check when R is a *field* and M and N are *finite-dimensional* that (5.4) is equivalent to equality of row and column rank for matrices. The injectivity of $\pi_L^\vee$ implies

$$\dim(\operatorname{im}(\pi_L^\vee)) = \dim((N/L(M))^\vee) = \dim(N/L(M)) = \dim N - \dim L(M). \tag{5.5}$$

Since $N^\vee/\ker(L^\vee) \cong L^\vee(N^\vee)$,

$$\dim(\ker(L^\vee)) = \dim N^\vee - \dim L^\vee(N^\vee) = \dim N - \dim L^\vee(N^\vee). \tag{5.6}$$

Because $\operatorname{im}(\pi_L^\vee) = \ker(L^\vee)$, comparing (5.5) and (5.6) shows $\dim L(M) = \dim L^\vee(N^\vee)$. That is, the images of L and $L^\vee$ have the same dimension when R is a field and M and N are finite-dimensional. Since L and $L^\vee$ can be represented by transposed matrices using a suitable choice of bases, these two dimensions are the column rank and row rank of any matrix representation for L. We can think of (5.4) as the generalization to any module of equality of row rank and column rank for matrices over a field.

6. Pairings

In Euclidean space, we tend to think of $\mathbf{R}^n$ as its own dual space: every functional φ on $\mathbf{R}^n$ is dotting with a fixed vector: $\varphi(w) = v \cdot w$ for some $v \in \mathbf{R}^n$. In Example 2.8 we saw that the dual R-module to the ideal $(3, 1 + \sqrt{-14})$ in $\mathbf{Z}[\sqrt{-14}]$ can be viewed as the ideal $(3, 1 - \sqrt{-14})$. There are many other settings where one module can play the role of the dual to another module even though it is not at first defined as the abstract dual module. This is formalized by the notion of a pairing of R-modules, as follows.

Definition 6.1. A *pairing* of M and N is a bilinear function $\langle \cdot, \cdot \rangle\colon M \times N \to R$: it is linear in each component when the other one is fixed. That is,

- $\langle m + m', n\rangle = \langle m, n\rangle + \langle m', n\rangle, \langle rm, n\rangle = r\langle m, n\rangle$,
- $\langle m, n + n'\rangle = \langle m, n\rangle + \langle m, n'\rangle, \langle m, rn\rangle = \langle m, rn\rangle$.

Example 6.2. The dot product is a pairing $\mathbf{R}^n \times \mathbf{R}^n \to \mathbf{R}$.

Example 6.3. For any commutative ring R, the dot product is a pairing $R^n \times R^n \to R$. In particular, taking $n = 1$, multiplication on R is a pairing.

Example 6.4. There are two "natural" pairings $\mathrm{M}_n(R) \times \mathrm{M}_n(R) \to R$: $\langle A, B\rangle = \mathrm{Tr}(AB)$ and $\langle A, B\rangle = \mathrm{Tr}(AB^\top)$. (What about $\mathrm{Tr}(A^\top B)$? Since $(A^\top B)^\top = B^\top A$, $\mathrm{Tr}(A^\top B) = \mathrm{Tr}(B^\top A) = \mathrm{Tr}(AB^\top)$, so we get nothing new.)

Example 6.5. When $R = \mathbf{Z}[\sqrt{-14}]$, Example 2.8 gave us the pairing $(3, 1+\sqrt{-14}) \times (3, 1-\sqrt{-14}) \to R$ where $\langle x, y\rangle = \frac{xy}{3}$.

Example 6.6. Let $\langle \cdot, \cdot \rangle \colon \mathbf{R}[X] \times \mathbf{R}[X] \to \mathbf{R}$ by $\langle f, g\rangle = f(0)g(0)$. This has the feature that $\langle X, g\rangle = 0$ for all g although $X \neq 0$ in $\mathbf{R}[X]$. More generally, $\langle f, g\rangle = 0$ for all g when $X | f$.

Example 6.7. Let $\langle \cdot, \cdot \rangle \colon M \times M^\vee \to R$ by $\langle m, \varphi\rangle = \varphi(m)$. This is the standard pairing between any module and its dual module.

Example 6.8. In analysis, if $p, q > 1$ satisfy $1/p + 1/q = 1$ then there is a pairing $L^p[0, 1] \times L^q[0, 1] \to \mathbf{R}$ given by $\langle f, g\rangle = \int_0^1 f(x)g(x)\,\mathrm{d}x$

Example 6.9. In topology, if X is a manifold then there is a pairing $\Omega^1(X) \times H^1(X, \mathbf{R}) \to \mathbf{R}$ of differential forms and cohomology classes given by integration: $\langle \omega, \gamma\rangle = \int_\gamma \omega$.

When we have a pairing $\langle \cdot, \cdot \rangle \colon M \times N \to R$, we can use it to let M and N behave as "part" of the dual to the other module: for each $m \in M$, $n \mapsto \langle m, n\rangle$ is a functional on N. Similarly, for each $n \in N$, $m \mapsto \langle m, n\rangle$ is a functional on M. Of course, if the pairing is badly behaved we could have $\langle m, n\rangle = 0$ for all n with $m \neq 0$. See Example 6.6, for instance. In any case, let's record the exact connection between pairings and maps into dual modules, which is the reason that pairings are mathematically interesting.

Definition 6.10. For R-modules M, N, the bilinear maps $M \times N \to R$ are denoted $\mathrm{Bil}_R(M, N; R)$.

Under the usual addition and scaling of bilinear maps, $\mathrm{Bil}_R(M, N; R)$ is an R-module.

Theorem 6.11. *For R-modules M, N, the modules*

$$\mathrm{Bil}_R(M, N; R), \ \mathrm{Hom}_R(M, N^\vee), \ \textit{and} \ \mathrm{Hom}_R(N, M^\vee)$$

are isomorphic.

Proof. To show $\mathrm{Hom}_R(M, N^\vee) \cong \mathrm{Bil}_R(M, N; R)$, let $L \colon M \to N^\vee$ be linear. Then to each $m \in M$ we have a linear map $L(m) \in N^\vee$. So from L we get a pairing $M \times N \to R$ by $\langle m, n\rangle_L = L(m)(n)$. To see $\langle \cdot, \cdot \rangle_L$ is a pairing, note it is linear in m with n fixed since L is linear and it is linear in n with m fixed since $L(m)$ is linear. Easily $\langle \cdot, \cdot \rangle_{L+L'} = \langle \cdot, \cdot \rangle_L + \langle \cdot, \cdot \rangle_{L'}$ and $\langle \cdot, \cdot \rangle_{rL} = r\langle \cdot, \cdot \rangle_L$. Thus $L \mapsto \langle \cdot, \cdot \rangle_L$ is a linear map from $\mathrm{Hom}_R(M, N^\vee)$ to $\mathrm{Bil}_R(M, N; R)$. Conversely, given a pairing $\langle \cdot, \cdot \rangle \colon M \times N \to R$, define $L \colon M \to N^\vee$ by $L(m) = \langle m, -\rangle$. The reason $L(m)$ is in $N^\vee$ is because $\langle \cdot, \cdot \rangle$ is linear in its second component with the first component fixed. The reason L is linear is because $\langle \cdot, \cdot \rangle$ is linear in its first component with the second component fixed. We have described how to pass from $\mathrm{Hom}_R(M, N^\vee)$ to $\mathrm{Bil}_R(M, N; R)$ and conversely. The reader can check these correspondences are inverses to each other, so we get R-module isomorphisms.

In the same way there is an isomorphism $\mathrm{Hom}_R(N, M^\vee) \to \mathrm{Bil}_R(M, N; R)$: if $L \in \mathrm{Hom}_R(N, M^\vee)$ let $\langle m, n\rangle_L = L(n)(m)$. This defines a pairing $M \times N \to R$, and if we are given a pairing $\langle \cdot, \cdot \rangle \colon M \times N \to R$ we get an element $L \in \mathrm{Hom}_R(N, M^\vee)$ by $L(n) = \langle -, n\rangle$.

Finally, we can write down an isomorphism between

$$\mathrm{Hom}_R(M, N^\vee) \cong \mathrm{Hom}_R(N, M^\vee)$$

as follows. If $L \in \mathrm{Hom}_R(M, N^\vee)$ let $L' \in \mathrm{Hom}_R(N, M^\vee)$ by $L'(n)(m) = L(m)(n)$. It is left to the reader to write down a correspondence in the other direction and check it is inverse to the one we have described.[1] □

Remark 6.12. For finite free R-modules M and N, the correspondence in Theorem 6.11 between $\mathrm{Hom}_R(M, N^\vee)$ and $\mathrm{Hom}_R(N, M^\vee)$ is precisely the dual map isomorphism in Corollary 5.10, with $N^\vee$ here in the role of N there and $N^{\vee\vee}$ identified with N by double duality.

A pairing $M \times N \to R$ lets us use M to parametrize a piece of $N^\vee$ and N to parametrize a piece of $M^\vee$. Of course, as Example 6.6 shows us, some pairings may make different elements of M behave like the same element of $N^\vee$ (*e.g.*, a non-zero element of M might pair with every element of N to the value 0). The pairings which let us identify M and N with each other's full dual module are the most important ones.

Definition 6.13. A pairing $\langle \cdot, \cdot \rangle \colon M \times N \to R$ is called *perfect* if the induced linear maps $M \to N^\vee$ and $N \to M^\vee$ are both isomorphisms.

The pairings at the start of this section are perfect except Examples 6.6, 6.7, and 6.8. Example 6.7 is perfect if and only if the natural map $M \to M^{\vee\vee}$ is an isomorphism (M is reflexive). Example 6.8 is perfect if we use the *continuous* dual space (continuous linear maps to $\mathbf{R}$).

When M and N are *finite free* modules of the *same* rank, to check a pairing $M \times N \to R$ is perfect it suffices to check that the induced linear map $M \to N^\vee$ is an isomorphism; the induced linear map $N \to M^\vee$ will then *automatically* be an isomorphism since it is just the dual to the linear map $M \to N^\vee$ (Remark 6.12). If R is a field and M and N have the *same* finite dimension, then a pairing $\langle \cdot, \cdot \rangle \colon M \times N \to R$ is perfect if and only if the induced map $M \to N^\vee$ is *injective* (that is, $\langle m, n \rangle = 0$ for all n only when $m = 0$ or equivalently when $m \neq 0$ there is an n such that $\langle m, n \rangle \neq 0$) since an injective linear map of vector spaces with the same dimension is an isomorphism. Back in the non-field case, an injective linear map of free modules with the same rank need not be an isomorphism (think of doubling $\mathbf{Z} \to \mathbf{Z}$), so to check a pairing $M \times N \to R$ of *finite free* modules M and N with the same rank is perfect when R is *not* a field, it is inadequate to check the induced linear map $M \to N^\vee$ is one-to-one (although it suffices to check it is onto, by a theorem of Vasconcelos).

The notion of a perfect pairing of modules M and N is *stronger* than an identification of just one of these modules with the dual of the other module: it identifies each module as the dual of the other: $M \cong N^\vee$ and $N \cong M^\vee$ (both isomorphisms coming from the same perfect pairing $M \times N \to R$). In particular, the existence of a perfect pairing of M and N forces the natural map $M \to M^{\vee\vee}$ to be an isomorphism. So only reflexive modules could be candidates for being part of a perfect pairing. In any event, the point of pairings is to systematize the idea that we can sometimes describe the dual of one module in terms of another known module (instead of working with an abstract dual module) and *vice versa*.

References

[1] J. Rotman, "Advanced Modern Algebra," Prentice-Hall, Upper Saddle River, NJ, 2002.

[1] If we use bilinear pairings $M \times N \to P$ into an R-module P, Theorem 6.11 extends to isomorphisms among the Hom-modules $\mathrm{Bil}_R(M, N; P)$, $\mathrm{Hom}_R(M, \mathrm{Hom}_R(N, P))$, and $\mathrm{Hom}_R(N, \mathrm{Hom}_R(M, P))$.