Propaganda. Let us briefly set up some notation. Let $\mathscr{C}$ be a category and $(I, \leq)$ a preordered set (or even better, some small category).

(a) We shall denote $R-\mathbf{mod}$ by $\mathbf{Rmod}$ (and right R-modules by $\mathbf{modR}$).
(b) We shall denote $\varinjlim_{(I,\leq)} \mathfrak{X}$ by $\operatorname{colim}\mathfrak{X}$. The mantra is "direct limits are colimits and inverse limits are limits."
(c) By a constant $(I, \leq)$ or I-shaped diagram in $\mathscr{C}$, we shall mean the functor $F\colon (I,\leq) \to \mathscr{C}$ which sends every object to one object of $\mathscr{C}$ and every arrow to the identity arrow. We often denote this by $\underline{X}$ where X is this object in $\mathscr{C}$.

Remark. To show a functor F is exact, it is enough to show it preserves exact sequences $X \to Y \to Z$ for any such exact sequence. This implies (and is equivalent to) preserving short exact sequences since this then says that $0 \to FA \to FB$ is exact, $FA \to FB \to FC$ is exact and $FB \to FC \to 0$ is exact, so that $0 \to FA \to FB \to FC \to 0$ is exact. The converse is obtained by sticking kernels and cokernels in the right places. It will be easier to deal with these shorter sequences.

Exercise (#1). *Let $\mathscr{C}$ be a category and $\mathscr{A}$ be an abelian category. Then $\mathscr{A}^{\mathscr{C}}$ (and similarly $\mathscr{A}^{\mathscr{C}^{\text{op}}}$) is again an abelian category.*

Proof. The point is this: As with most interesting function spaces, we tend to pick up important properties of the codomain. For instance, if X and Y are normed linear spaces with Y complete, then $L(X,Y)$ picks up the operator norm and is complete with respect to it. In particular, we are allowed to compute things "pointwise."

Denote $\mathbf{F} = \mathscr{A}^{\mathscr{C}}$. Let $F, G \in \mathscr{A}^{\mathscr{C}}$ and let $\eta, \tau \in \operatorname{Hom}_{\mathbf{F}}(F,G)$. Then $\eta + \tau \overset{\text{def}}{=} (\eta_X + \tau_X)_{X\in\mathscr{C}}$ where $\eta_X, \tau_X \in \operatorname{Hom}_{\mathscr{A}}(FX, GX)$ and, therefore, we are permitted to add them in $\operatorname{Hom}_{\mathscr{A}}(FX,GX)$ since $\mathscr{A}$ is additive. We assert that this is a natural transformation $F \Rightarrow G$. Since composition is bilinear on components since $\mathscr{A}$ is additive, if $f \in \mathscr{C}$ is an arrow, then $Ff \circ (\eta+\tau) = Ff\circ\eta + Ff\circ\tau = \eta\circ Gf + \tau\circ Gf$ from naturality of η and τ and bilinearity. So this is a natural transformation. Having defined addition ***pointwise*** (or, if you prefer to think about it this way, ***componentwise***) on the hom-sets, it is immediate that the additive identity $\mathbf{0} \in \operatorname{Hom}_{\mathbf{F}}(F,G)$ is the natural transformation which is the additive identity 0 pointwise on each component. Moreover, given $\eta \in \operatorname{Hom}_{\mathbf{F}}(F,G)$, $-\eta \overset{\text{def}}{=} (-\eta_X)_{X\in\mathscr{C}}$ is the inverse of η on each component and therefore clearly also for η. Since we have defined addition *pointwise*, it is obviously associative because it picks this up from associativity down in $\mathscr{A}$. This proves that for every $F, G \in \mathbf{F}$, $\operatorname{Hom}_{\mathbf{F}}(F,G)$ is a group under the defined operation. It is so obviously abelian that it scarcely merits mentioning.

We claim that composition is $\mathbf{Z}$-bilinear in $\mathbf{F}$ under our addition. Let $\eta,\tau\colon F \Rightarrow G$, and $\sigma\colon G \Rightarrow H$ and consider $\sigma\circ(n\eta + m\tau)$ where $n, m \in \mathbf{Z}$. Since we have defined everything *pointwise*, it suffices to check bilinearity point-by-point. Consider the component at $X \in \mathscr{C}$. Then we have $\sigma_X \circ (n\eta_X + m\tau_X) \in \operatorname{Hom}_{\mathscr{A}}(FX, HX)$—hence, since $\mathscr{A}$ is abelian, $\sigma_X\circ(n\eta_X + m\tau_X) = n\sigma_X\circ\eta_X + m\sigma_X\circ\tau_X$. Thus, composition is $\mathbf{Z}$-bilinear.

We claim that there exists a zero-object in $\mathbf{F}$. This is obvious. Fix $0 = 0_{\mathscr{A}} \in \mathscr{A}$ a zero-object in $\mathscr{A}$—this exists because $\mathscr{A}$ is additive. The functor $\mathcal{O}\colon \mathscr{C} \to \mathscr{A}$ which sends each object $X \in \mathscr{C}$ to 0 and each arrow $f\colon X \to Y$ to $0 \in \operatorname{Hom}_{\mathscr{A}}(0,0)$. Then there is clearly only one natural transformation in $\operatorname{Hom}_{\mathbf{F}}(\mathcal{O}, F)$ and $\operatorname{Hom}_{\mathbf{F}}(F,\mathcal{O})$ for any functor $F \in \mathbf{F}$—the trivial one.

Let $F_1, F_2 \in \mathbf{F}$. Define $F_1 \oplus F_2$ for each $X \xrightarrow{f} Y$ by $(F_1\oplus F_2)(f) = F_1(f)\oplus F_2(f)$ where $F_1(f)\oplus F_2(f)\colon F_1X\oplus F_2X \to F_1Y \oplus F_2Y$ is defined by the universal property of coproduct. Note that this exists since $\mathscr{A}$ is additive. Define a natural transformation $\eta^i\colon F_i \Rightarrow F_1\oplus F_2$ whose component at each $X\in\mathscr{C}$ is the canonical map $\eta^i_X\colon F_iX \to F_1X\oplus F_2X$ guaranteed as part of the universal property of the coproduct. Note that the map $F_1f\oplus F_2f$ is the map induced by $\eta^i_Y \circ F_if$ for $i = 1, 2$. Hence, it follows that for all $X \xrightarrow{f} Y$, $\eta^i_Y\circ F_if = (F_1f\oplus F_2f)\circ\eta^i_X$, from the universal property. Hence, η^i really is a natural transformation for $i = 1,2$. Thus, it follows by a trivial induction that $\mathbf{F}$ has finite coproducts. Similarly for products, $\prod F_i$ is simply the pointwise product on the components. This proves that $\mathbf{F}$ is at least additive.

We shall show that $\mathbf{F}$ contains all cokernels—that $\mathbf{F}$ contains all kernels is analogous and the proof is omitted. For clarity of presentation, we omit $\Rightarrow$ for natural transformations. Let $\eta\colon F \to G$ be a natural transformation. For all $X \xrightarrow{f} Y \xrightarrow{g} Z$ in $\mathscr{C}$, we have TFCD:

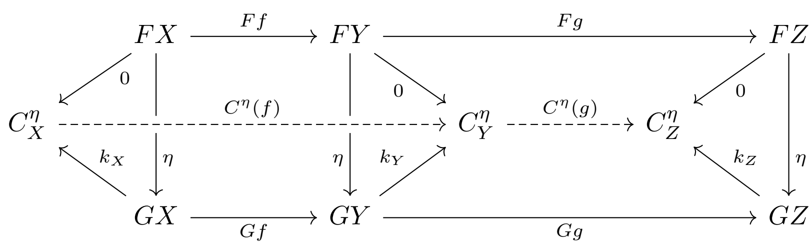

where each the dashed arrows are induced by the universal property of the cokernel (just follow the appropriate portions of the diagram) and where $k_X\colon GX \to C^\eta_X$ is the cokernel of the component η_X for each $X \in \mathscr{C}$. This exists because this is now taking place in the abelian category $\mathscr{A}$. In particular, uniqueness of such arrows forces $C^\eta(g \circ f) = C^\eta(g) \circ C^\eta(f)$. In other words, the assignment $f \mapsto C^\eta(f)$ is functorial. Hence, C^η is a functor (since objects correspond to identity arrows and conversely, we see this is enough). Let $k \colon G \to C^\eta$ be the natural transformation whose components are the k_X as above—that this is a natural transformation follows directly from the commutativity of the diagram above. We assert that $k \colon G \to C^\eta$ is the cokernel of η. Once again, since we have defined everything *pointwise*, it suffices to check point-by-point—towards this end, we note that each k_X is a cokernel so this follows immediately because we have contrived it to be so. However, we need to work this out in full detail, apparently, so let's go through the motions. Suppose we have TFCD in $\mathbf{F}$:

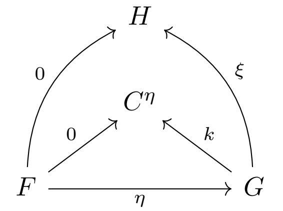

(1)

where 0 obviously denotes the trivial natural transformation. Opening up this picture on components (hence, why we don't really need to check) we have TFCD for each $X \xrightarrow{f} Y$ in $\mathscr{C}$:

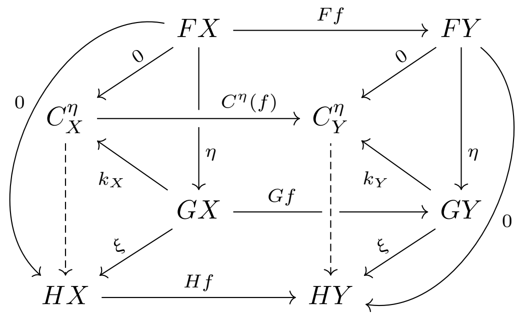

where the dashed arrows are induced by the universal property of the cokernel by way of $\xi \colon G \to H$ and $0 \colon F \to H$—hence, they are uniquely defined. They make every such diagram commute. It is therefore the case that the dashed arrows form the unique natural transformation fitting into (1) making the diagram commute and hence that we really do have a cokernel in $\mathbf{F}$.

With our point-by-point computation mantra, it is obvious that the induced maps $\operatorname{Coim} f \to \operatorname{Im} f$ are isomorphisms. Hence, $\mathscr{A}^{\mathscr{C}}$ is an abelian category. To see that $\mathscr{A}^{\mathscr{C}^{\mathrm{op}}}$ is an abelian category, note that we already showed that $\mathscr{A}^{\mathscr{C}}$ is an abelian category for any category $\mathscr{C}$. Hence, for all categories under consideration. ■

Exercise (#7). *Let J be a set. Then $\bigoplus_J, \prod_J \colon \mathbf{Rmod}^J \to \mathbf{Rmod}$ are exact functors.*

Proof. We view J as a discrete category—in this case, J has objects consisting of its elements and its only morphisms are identity morphisms. Note that **Rmod** contains all products and coproducts. By **Exercise 1**, $\mathbf{Rmod}^J$ is an abelian category.

We may as well note the following.

Lemma. *Kernels, cokernels and images are what we expect them to be in* **Rmod**.

Proof. Trivial. ■

Let $F_1, F_2, F_3 \in \mathbf{Rmod}^J$ be such that we have an exact sequence $F_1 \overset{\tau}{\Rightarrow} F_2 \overset{\varepsilon}{\Rightarrow} F_3$. Unraveling the definition, this amounts to an exact sequence in **Rmod** for each $j \in J$,

$$F_1(j) \xrightarrow{\tau_j} F_2(j) \xrightarrow{\varepsilon_j} F_3(j)$$

Let us first consider the case of $\bigoplus_J$. We know how to describe direct sums in **Rmod**—they can be realized as the direct sum of abelian groups with the component-wise R-action along with the inclusion maps of the direct summands into the direct sum at its component. In this case, $\bigoplus \varepsilon_j$ is the map which sends an element $(x_j)_{j \in J} \in \bigoplus_J F_2(j)$ to $(\varepsilon_j(x_j))$—as always with maps induced from direct sums, this is well-defined because only finitely many x_j are nonzero. It is also clearly an R-linear map. The kernel of the map is necessarily the set where the ε_j's simultaneously vanish. That is, $\operatorname{Ker} \bigoplus \varepsilon_j = \bigoplus \operatorname{Ker} \varepsilon_j$. It is also obvious that $\operatorname{Im} \bigoplus \tau_j = \bigoplus \operatorname{Im} \tau_j$ by elementary set-theoretic considerations. Since $\operatorname{Ker} \varepsilon_j = \operatorname{Im} \tau_j$ for each j, it follows immediately that $\bigoplus_J$ is exact.

The proof for the direct product is hardly different and is therefore omitted, as was suggested in recitation. The idea is the same as before. ■

Exercise (#8). *Let $(I, \leq)$ be a right-directed, preordered set. Then* $\operatorname{colim}\colon \mathbf{Rmod}^{(I,\leq)} \to \mathbf{Rmod}$ *is an exact functor.*

Proof. I don't think we've seen the existence of direct limits in **Rmod** yet, so we must justify their existence before proceeding. This isn't hard and we make it brief.

Lemma (Existence). *The colimit of any $(I, \leq)$-shaped diagram in* **Rmod** *exists.*

Proof. For $F\colon (I, \leq) \to \mathbf{Rmod}$, $\operatorname{colim} F$ is simply the quotient of $\bigoplus_I F$ by the submodule Q generated by all elements $(x_i) \in \bigoplus_I F$ of the form $(\ldots, 0, \ldots, 0, x_i, 0 \ldots, 0, -f_{ji}(x_i), 0, \ldots)$ where $f_{ji}\colon F(i) \to F(j)$ is as guaranteed by the directed system in **Rmod** obtained from the functor F (i.e., where $i \leq j$). The maps of $F(i)$ into this object are the natural ones: The composite of the inclusion and canonical projection $\iota_i\colon F(i) \hookrightarrow \bigoplus_I F \to \bigoplus_I F/Q$. Given a constant $(I, \leq)$-shaped diagram in **Rmod**, $\underline{Y}$, say, with R-linear maps $\kappa_j\colon F(j) \to Y$, the universal property of the direct sum yields an R-linear map $f\colon \bigoplus_I F \to Y$ such that $f|\, F(i) = \kappa_i$ in the obvious sense. In particular, since $\kappa_j(f_{ji}(x_i)) = \kappa_i(x_i)$, it must be that $Q \subseteq \operatorname{Ker} f$. Since f is constant on the fibers of the quotient map $p\colon \bigoplus_I F \to \bigoplus_I F/Q$, it factors through p uniquely (i.e., we obtain a unique R-linear map $\bigoplus_I F/Q \to Y$). It is clear that $f \circ \iota_i = \kappa_i$. Conversely, if we have an R-linear map $h\colon \bigoplus_I F/Q \to Y$ which is a morphism of constant $(I, \leq)$-shaped diagrams, then we obtain a map $h \circ p\colon \bigoplus_I F \to Y$ which, again as a result of the the universal property of the coproduct, induces a map $f\colon \bigoplus_I F \to Y$. Running the same argument through as above shows that f descends uniquely to an R-linear map $\tilde{h}\colon \bigoplus_I F/Q \to Y$ for which $\tilde{h} \circ p = h \circ p$. Since h is unique, $\tilde{h} = h$. This proves that $\bigoplus_I F/Q \approx \operatorname{colim} F$. ■

Consider $(I, \leq)$-shaped diagrams in **Rmod**, (F, f_{ji}), (F', f'_{ji}) and (F'', f''_{ji})—here we have denoted the compatible maps along with them—and suppose we are given an exact sequence $F \overset{\tau}{\Rightarrow} F' \overset{\varepsilon}{\Rightarrow} F''$. We already know that $\bigoplus$ is an exact functor from **Exercise 7**. We assert that we have TFCD for all $i, j \in I$ with $i \leq j$:

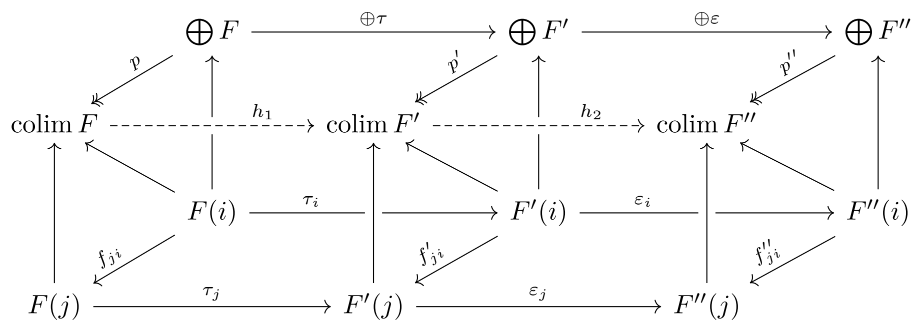

We have obtained this from our construction of the colimit, the universal property of the direct sum applied to the components τ (resp. ε) along with the canonical maps into the direct sum given by the universal property. The dashed horizontal arrows are the unique ones induced by the universal property of the colimit. Note that each p is the unique arrow guaranteed by the universal property of the coproduct which is induced by the maps into the colimit—this is because we constructed the colimit as a quotient of the direct sum. Moreover, in this diagram, the horizontal arrows in the topmost and both bottommost rows form an exact sequence as a direct consequence of the previous exercise and our assumptions, respectively. The unnamed arrows into the colimits are the obvious ones. Now, since the top row is exact, $0 = p'' \circ \oplus\varepsilon \circ \oplus\tau = h_2 \circ h_1 \circ p$. Since p surjects, this shows that $\operatorname{Im} h_1 \subseteq \operatorname{Ker} h_2$.

Conversely, suppose $x \in \operatorname{Ker} h_2$. Then x is a finite sum of elements in the images of $F'(i) \to \operatorname{colim} F'$ varying over i. Pick representatives of this sum, say $x_{i_n} \in F'(i_n)$ for finitely many n, so that x is the sum of the images of the x_{i_n} in $\operatorname{colim} F'$. By right-directedness of I, there exists j which dominates all elements witnessing this sum. Let $y = \sum f'_{ji_n}(x_{i_n})$ so that $y \in F'(j)$. By R-linearity of the f'_{ji} and compatibility of the maps, the image of y in $\operatorname{colim} F'$ is x.

Now, $\varepsilon_j(y) \in \operatorname{Ker}(F''(j) \to \operatorname{colim} F'')$ by exactness of the bottom row. Identify $\varepsilon_j(y)$ with its image in $\bigoplus F''$. By definition of the colimit and p'', this means that $\varepsilon_j(y)$ is an R-linear combination of elements of the form $x_{i_n} - f''_{ji_n}(x_{i_n})$—note that we have implicitly lifted everything to $\bigoplus F''$ in order to make sense of this sum. Since $\varepsilon_j(y) \in F''(j)$, the sum $y = \sum_n x_{i_n} - f''_{ji_m}(x_{i_n})$ vanishes off of $F''(j)$. Now, by the same argument we used before, we can push this sum forward to obtain $y' \in F''(k)$ for which $y' \in \operatorname{Ker}(F''(k) \to \operatorname{colim} F'')$. But also

$$y' = \sum_n f''_{ki_n}(x_{i_n}) - f''_{kj}(f''_{ji_n}(x_{i_n})) = \sum_n f''_{ki_n}(x_{i_n}) - f''_{ki_n}(x_{i_n}) = 0.$$

We assert $0 = f''_{kj}(\varepsilon_j(y))$. Recall that $\varepsilon_j(y) = \sum_n x_{i_n} - f''_{ji_n}(x_{i_n})$ in the direct sum. It must be that terms not lying in the inclusion $F''(j) \to \bigoplus F''$ cancel, so it follows immediately from the centered equation above that $f''_{kj}(\varepsilon_j(y)) = 0$, since equality is unambiguous. By commutativity of the diagram, this implies that $\varepsilon_k(f'_{kj}(y)) = 0$ so that $f'_{kj}(y) \in \operatorname{Ker} \varepsilon_k$; hence, there exists $z \in F(k)$ such that $\tau_k(z) = f'_{kj}(y)$. But by commutativity of the diagram, this implies that $h_1([z]) \in \operatorname{Ker} h_2$, where $[z]$ is the image of z in $\operatorname{colim} F$, and furthermore that $h_1([z]) = x$. Hence, $\operatorname{Ker} h_2 = \operatorname{Im} h_1$. Hence, the functor is exact. ■